



Provided by the author(s) and University of Galway in accordance with publisher policies. Please cite the published version when available.

Title	Modularity, variation and variability in genetic representation
Author(s)	Hill, James
Publication Date	2017-12-05
Publisher	NUI Galway
Item record	http://hdl.handle.net/10379/10029

Downloaded 2024-05-18T18:42:25Z

Some rights reserved. For more information, please see the item record link above.





Modularity, Variation and Variability in Genetic Representation

A dissertation presented

by

James Hill

to

The College of Engineering & Informatics

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in the subject of

Computer Science

Supervisor: Dr. Colm O'Riordan

National University of Ireland Galway

December 2017

©2017 - James Hill

All rights reserved.

Thesis Supervisor

Author

Dr. Colm O’Riordan

James Hill

Modularity, Variation and Variability in Genetic Representation

Abstract

Nature uses a complex genotype-phenotype map to advance a relatively simple genotype space variational topology to an extremely complex phenotypic variational topology. This dissertation introduces a modular, fixed non-trivial, multi-layered genotype-phenotype map, incorporating an interpretation of the biological processes of *transcription* and *translation* into the representation of a Genetic Algorithm (GA), thereby introducing a more flexible phenotypic structure and increasing connectivity. A series of experiments are conducted to examine the impact of the representation on variation and variability, with results indicating that through the inclusion of the proposed multi-layered mapping, only a small fraction of genotypic mutations are adaptive. Results also suggest that slowing genetic drift has an important part to play in the evolution of a population, allowing the representation to foster the creation of highly connected neutral networks, altering genetic drift and changing the evolutionary trajectory by inducing variability.

Contents

Title Page	i
Abstract	iii
Table of Contents	iv
List of Figures	vii
List of Tables	ix
Declaration	x
Citations to Previously Published Work	xi
Acknowledgments	xiii
1 Introduction	1
1.1 Genetic Algorithms	1
1.2 Biological Concepts	5
1.3 Modularity & GA Representation	8
1.4 Open Research Questions	10
1.5 Dissertation Goals and Hypothesis	11
1.5.1 Motivation	14
1.5.2 Methodology	14
1.5.3 Contribution	15
1.6 Thesis Structure	15
1.7 Chapter Summary	16
2 Genetic Algorithm	
Fundamentals & Performance Measurement	17
2.1 Introduction	17
2.2 Genetic Algorithms - Above the Surface	18
2.3 Genetic Algorithms - Below the Surface	21
2.4 Exploring the Search Space	25
2.5 The Fitness Landscape	28
2.6 Measuring GA Performance	31
2.6.1 Static Landscapes	32
2.6.2 Dynamic Landscapes	42

2.7	Chapter Summary	47
3	Neutrality	49
3.1	Introduction	49
3.2	Neutral Theory & Neutrality	50
3.3	Synthetic Neutrality	54
3.3.1	Types & Properties	54
3.3.2	Neutrality in Artificial Systems	56
3.3.3	Neutrality in GAs	61
3.4	Chapter Summary	72
4	Multi-Layered GA (MGA) Design	74
4.1	Introduction	74
4.2	Motivating the MGA	75
4.2.1	Phenotypic Variability & Neighbourhood Equivalence	77
4.3	MGA Design	78
4.4	MGA Representation	81
4.4.1	Transcription Phase	81
4.4.2	Translation Phase	83
4.4.3	Missense Mutation Operator	88
4.4.4	MGA Arity	89
4.5	MGA Mapping Overview	92
4.6	Chapter Summary	95
5	Examining the Efficacy of a Multi-Layered GA	97
5.1	Introduction	97
5.2	Efficacy Experiments Overview	98
5.3	Statistical Significance	99
5.4	Test 1 - Deceptive Landscapes	99
5.4.1	Test 1 Experiment Results	100
5.5	Test 2 - De Jong's Test Suite	107
5.5.1	Test 2 Experiment Results	107
5.6	Chapter Summary	114
6	Examining Population Variation in a Multi-layered GA.	117
6.1	Introduction	117
6.2	Variation Experiments Overview	118
6.3	Test 3 - Inducing Phenotypic Variability	119
6.3.1	Test 3 Experiment Results	120
6.4	Test 4 - Altering Dimensionality	125
6.4.1	Test 4 Experiment Results	126
6.5	Chapter Summary	137

7	Examining the Arity of a Multi-layered GA	140
7.1	Introduction	140
7.2	Tuning the MGA Arity	141
7.3	Arity Experiments Overview	143
7.4	Test 5 - Altering the Arity	144
7.4.1	Test 5 Experiment Results	145
7.5	Chapter Summary	162
8	Examining the Population Trajectory & Heterogeneity of a Multi-layered GA	166
8.1	Introduction	166
8.2	4-bit Fully Deceptive Landscape	168
8.3	Experiments Overview	171
8.3.1	Population Trajectory Experiments Overview	171
8.3.2	Heterogeneity Experiments Overview	171
8.4	Test 6 - Evolutionary Trajectory	172
8.4.1	Neutrality & Population Evolution	173
8.4.2	Neutrality & Variation	178
8.4.3	Neutrality & Variability	182
8.5	Test 7 - Heterogeneity	186
8.5.1	Test 7 Experiments	188
8.5.2	Neutral Networks & Neighbourhood Equivalence	189
8.5.3	Selection Mechanisms	192
8.5.4	Selection Noise	196
8.5.5	Selection Pressure	198
8.5.6	Operator Disruption	204
8.6	Chapter Summary	206
9	Discussion & Conclusion	211
9.1	Introduction	211
9.2	Discussion	216
9.2.1	Efficacy Results	217
9.2.2	Variation Results	220
9.2.3	Variability Results	225
9.3	Conclusion	230
9.4	Future Work	233
	Bibliography	235

List of Figures

2.1	The Sphere Model.	39
2.2	Rosenbrock’s Function.	39
2.3	Step Function.	40
2.4	Quadratic Function.	41
2.5	Shekel’s Foxholes.	42
2.6	4-bit Changing Deceptive Landscape	46
4.1	Overview of Multi-layered GA (MGA)	79
4.2	Non-Redundant & Redundant Connectivity	86
4.3	4-bit MGA Representation Mapping	90
4.4	6-bit MGA Representation Mapping	91
4.5	8-bit MGA Representation Mapping	91
5.1	Off-line Performance 3-bit Deceptive Problem	101
5.2	Off-line Performance 10 3-bit Deceptive Problem	103
5.3	Sub-Functions Discovered 10 3-bit Deceptive Problem	103
5.4	Off-line Performance 10 3-bit Changing Landscape	105
5.5	Sub-functions Discovered 10 3-bit Changing Landscape	106
5.6	The Sphere Model - SGA & MGA.	108
5.7	Rosenbrock’s Saddle - SGA & MGA	110
5.8	Step Function - SGA & MGA	111
5.9	Quadratic Function with Noise - SGA & MGA	112
5.10	Shekel’s Foxholes - SGA & MGA	113
6.1	Genotype/Phenotype Diversity 3-bit Deceptive Problem	121
6.2	Genotype/Phenotype Diversity 10 3-bit Deceptive Problem	123
6.3	Genotype/Phenotype Diversity 10 3-bit Changing Landscape	125
6.4	Sphere Model Diversity $n = 3$	127
6.5	Sphere Model Diversity $n = 15$	129
6.6	Sphere Model Diversity $n = 30$	131
6.7	Sphere Model Changing Environment Diversity $n = 3$	133

6.8	Sphere Model Changing Environment Diversity $n = 15$	135
6.9	Sphere Model Changing Environment Diversity $n = 30$	136
7.1	30-bit OneMax Static Landscape	146
7.2	Diversity 30-bit OneMax Static Landscape	147
7.3	90-bit One Max Static Landscape	149
7.4	Diversity 90-bit OneMax Static Landscape	150
7.5	10 3-bit Deceptive Changing Environment	152
7.6	Diversity 10 3-bit Deceptive Changing Environment	153
7.7	30 3-bit Deceptive Changing Environment	155
7.8	Diversity 30 3-bit Deceptive Changing Environment	157
7.9	No Transcription Changing Environment	159
7.10	No Missense Mutation Changing Environment	160
7.11	Increased Mutation Changing Environment	161
8.1	4-bit Deceptive Landscapes & Associated Heat Maps	169
8.2	SGA & MGA On-line/Off-line Performance Statistics	170
8.3	SGA & MGA Generation 0	173
8.4	SGA & MGA Generation 50	175
8.5	SGA & MGA Generation 51	176
8.6	SGA & MGA Generation 200	177
8.7	SGA & MGA Diversity - Generation 0	179
8.8	SGA & MGA Diversity - Generation 50	180
8.9	SGA & MGA Diversity - Generation 200	181
8.10	Single-bit Mutation SGA	183
8.11	Single-bit Mutation MGA	184
8.12	Single-bit & Missense Mutated Individual	185
8.13	Neutral Networks	190
8.14	Neutral Networks - Identical Genotype	191
8.15	SGA (Roulette Wheel) & MGA On-line/Off-line	193
8.16	SGA (Tournament) & MGA On-line/Off-line	193
8.17	SGA (Ranking) & MGA On-line/Off-line	194
8.18	SGA & MGA Convergence Rate	195
8.19	Selection Noise On-line/Off-line	197
8.20	Selection Noise Convergence Rate	198
8.21	Selection Pressure On-line/Off-line	200
8.22	Selection Pressure Convergence Rate	201
8.23	Niching On-line/Off-line	203
8.24	Niching Techniques Convergence Rate	204
8.25	Operator Disruption (Crossover) On-line/Off-line	204
8.26	Operator Disruption (Mutation) On-line/Off-line	205
8.27	Operator Disruption Convergence Rate	206

List of Tables

2.1	Deceptive Order-3 Problem - Fitness Values	35
2.2	Order-3 Deceptive Problem & Reversed Order-3 Problem	45
2.3	Deceptive & Reversed Deceptive Order-4 Problem - Fitness Values	47
4.1	Template, Coding & RNA Mapping	82
4.2	Binary Representation for Phenets '0' & '1'	84
4.3	Character Representation for Phenets '0' & '1'	84
4.4	MGA 6-3-1 Dynamic Activation Pairings	85
4.5	Hyperplane Sampling	87
4.6	Hyperplane Sampling	88
4.7	Missense Mutation Mapping	89
4.8	Genotype-Phenotype Mapping Example	92
4.9	Messense Mutation Mapping Example	93
6.1	Sphere Model $n = 3$	127
6.2	Sphere Model $n = 15$	129
6.3	Sphere Model $n = 30$	131
6.4	Sphere Model Changing Landscape $n = 3$	133
6.5	Sphere Model Changing Environment $n = 15$	134
6.6	Sphere Model Changing Environment $n = 30$	136
8.1	Off-Line & On-Line Selection Comparison	197
8.2	Off-Line Scaling Comparison	201

Declaration

I declare that the work contained in this thesis, which I submit to the National University of Ireland Galway in consideration of the award of a Doctor of Philosophy, is all of my own work apart from cited references.

Citations to Previously Published Work

Large portions of Chapters 3, 4, 5, 6, 7 and 8 have appeared in the following papers:

Solving Fully Deceptive Problems in Changing Environments.

Séamus Hill and Colm O’Riordan,

Proceedings of the 2010 Artificial Intelligence Cognitive Studies (AICS 2010),

National University of Ireland, Galway, 2010.

A Genetic Algorithm with a Multi-layered Genotype-Phenotype Mapping.

Séamus Hill and Colm O’Riordan,

Proceedings of the International Conference on Evolutionary Computation Theory and Applications (ECTA 2010),

Valencia, Spain, 2010.

Examining the use of a Non-Trivial Fixed Genotype-Phenotype Mapping in Genetic Algorithms to Induce Phenotypic Variability over Deceptive Uncertain Landscapes.

Séamus Hill and Colm O’Riordan,

Proceedings of the 2011 Congress of Evolutionary Computation (CEC 2011),

New Orleans, USA, 2011.

Neutrality through Transcription & Translation in Genetic Algorithm Representation.

Séamus Hill and Colm O’Riordan,

Proceedings of the 3rd International Conference on Evolutionary Computation Theory and Applications (ECTA 2012),

Barcelona, Spain, 2012.

Analysing the Impact of Dimensionality & Diversity in a Multi-Layered Genotype-Phenotype Mapped GA.

Séamus Hill and Colm O’Riordan,

Proceedings of the 2013 Congress of Evolutionary Computation (CEC 2013),

Cancun, Mexico, 2013.

Altering the Granularity of Neutrality in a Multi-Layered Genetic Algorithm.

Séamus Hill and Colm O’Riordan,

Proceedings of the 6th International Conference on Evolutionary Computation Theory and Applications (ECTA 2014),

Rome, Italy, 2014.

Examining the Impact of Neutral Theory on Genetic Algorithm Population Evolution.

Séamus Hill and Colm O’Riordan,

Proceedings of the 7th International Conference on Evolutionary Computation Theory and Applications (ECTA 2015),

Lisbon, Portugal, 2015.

Diversifying Techniques & Neutrality in Genetic Algorithms.

Séamus Hill and Colm O’Riordan,

Proceedings of the 8th International Conference on Evolutionary Computation Theory and Applications (ECTA 2016),

Porto, Portugal, 2016.

Acknowledgments

Ah Colm how you have suffered with this baby... Thank you so much!

I would also like to thank the external examiners Prof. Pier Luca Lanzi from the Politecnico di Milano and Prof. Juan Julian Merelo from the University of Granada; the internal examiner Dr. Conn Mulvihill, for their kind words and advice.

Chapter 1

Introduction

1.1 Genetic Algorithms

Evolutionary Algorithms (EAs), as a problem solving paradigm, attempt to simulate the process of evolution as found in nature in a computer environment [8, 162, 175, 197]. EAs are stochastic algorithms which use natural phenomena such as survival of the fittest and the inheritance of genetic material. Evolution can be characterised using various levels of hierarchy, such as species, individuals, chromosomes and genes [56]. Genetic Algorithms (GAs) [95, 67] are a type of EA and are based on natural occurring phenomenon such as reproduction, recombination, mutation and selection. GAs are search algorithms based on the Darwinian principle of survival of the fittest [39]. An initial population of individuals is created, each representing a possible solution to a given problem. The individuals in the population are allocated a fitness score based on their suitability towards the environment in which they exist [161]. The individuals are then subjected to environmental pressure and based on

their fitness scores, natural selection takes place, with fitter candidates being offered the chance to seed the next generation through the use of operators such as recombination and mutation. In this way a new population is created and represents a new generation. The two processes at work here are first, the use of operators such as mutation and recombination to create variation within the population and second, selection, which encourages quality. Ultimately the effectiveness of any EA is determined by the relationship between the landscape and the operations used to create new individuals [3].

The search process of a GA can be broken into two processes: *exploration*, where the new generation of individuals test unexplored areas of the search space, and *exploitation*, where the search focuses on areas where good solutions reside. Maintaining a balance between exploration and exploitation is critical for a GA to succeed in searching the search space. As the search progresses, GAs map from the *genotype* or the representation of the individual in the population, to the *phenotype*, which can be viewed as a manifestation of a possible solution to the particular problem. This leads to a situation where the search moves from the *genotype space* to the *phenotype space*. The relationship between the genotype and the phenotype is determined by the genotype-phenotype mapping. In standard GAs each phenotype is represented by a distinct genotype.

GAs typically tend to use a fixed one-to-one genotype-phenotype map; implying that each genotype relates to a specific phenotype [169]. This means that if evolution is based on a genotype-phenotype map where each phenotype is represented by a particular genotype, you arrive at a point where there is a static fitness value

associated with every genotype in the population, then a particular genotype which represents the local optimum would have to accept a lower fitness value to escape and begin searching for a fitter point in the landscape. This is because there is no neutrality in the representation. With a neutral representation each phenotype could be represented by a number of different genotypes.

However, if the genotype-phenotype map were to include neutrality then there is the possibility of continuing the search by *drifting along neutral networks in the genotype space* [169]. In other words we have, by increasing neutrality in the genotype-phenotype map, changed the genotypic neighbourhood, as there are a neighbourhood of genotypes that map to the same phenotype, so even when an individual becomes trapped they can continue the search without the loss of fitness. Therefore, by having a many-to-one genotype-phenotype map, we introduce these *phenotypically neutral* neighbours [169], which potentially, allow a passage through the genotype space without losing fitness when trapped.

The phenotype space can also be viewed as the phenotypic topological space which can also contain a set of neighbourhoods. However, the difficulty here is that phenotypes, by their nature, are not altered by physical processes such as mutation. To overcome this, neighbourhoods in the phenotypic topological space can be classified by closeness in terms of a continuous evolutionary path. Therefore, by examining the frequency by which a mutation on a particular genotype yields a phenotype, we can define a notion of closeness in the phenotypic topological space [59]. Neutrality and neutral genetic drift in the mapping process is important in defining the phenotypic topology, and the variation in phenotypic topology. Because of the presence of

neutrality, phenotypes can be used to mean similarity between genotypes. Thus, the closeness of phenotype neighbourhoods depends on neighbourhood closeness among genotype neighbourhoods [60]. The genotype-phenotype map generates a phenotypic neighbourhood, which very often differs from the idea of proximity which exists between phenotypes when we consider them as a whole [59]. The relationship between the genes and the phenes (phenotypic traits which have an influence on a phenotype's overall fitness) is controlled by the genotype-phenotype map [1].

GAs have a hierarchy, similar in concept, to those found in nature and *chromosomes* and *genes* are the elements involved in simulating evolution. But, due to the simplicity of the representation, attempts to make the algorithms *biologically accurate* are not possible due to the level of abstraction inherent in the elements of a GA [56]. Because of this, the operators used in association with GAs are basic interpretations of those found in nature. However, the concepts observed in nature, prove useful in providing GA practitioners with ideas and analogies to interpret and incorporate, at an abstract level, into GAs in an attempt to improve performance [56].

As time progresses, knowledge of the concepts found in nature changes. A question as to which force is dominant in relation to evolution has existed for many years. Is natural selection the primary force that shapes evolution or is it the variation of genetic material? Darwin [39] argued for selection, although without a knowledge of genetics, while Sewall Wright [203] noticed that many random changes in the frequency of alleles occurring in a population were not related to selection. Wright's observations indicated that this *genetic drift*, meaning a change in a population's gene frequency resulting from a random variation in the distribution of genes from

one generation to the next, was an important component in the evolutionary process. *Neutral theory* as proposed by Kimura [111], offered an alternative to the Darwinian view and states that the mutations involved in the evolutionary process are neither advantageous nor disadvantageous to the survival of an individual, and that most adaptive mutations are not due to selection but rather to random genetic drift. However, Kimura [112] pointed out that although natural selection does play a role in adaptive evolution, only a tiny fraction of DNA changes are adaptive. The vast bulk of mutations are phenotypically silent. Nei [139] argues that all evolution must be non-Darwinian and as phenotypic characters are under the control of DNA sequences, molecular and phenotypic evolution must take place under a similar non-Darwinian approach. Nei also argues that the majority of *morphological evolution* is as a result of neutral or nearly neutral mutation and therefore mutation can be considered the main force behind evolution at both the molecular level and the phenotypic level [139].

1.2 Biological Concepts

In biological studies, Mendel discovered that nature holds all of the genetic information relating to an individual in pairwise alleles [153]. This genetic information determines the various physical characteristics that go to make up a particular individual. Later it was discovered that genes are made up from a long molecule known as DNA, which is copied and inherited from generation to generation. The genetic information is made up of a double string composed of four nucleotides (A, C, G and U) and line up in a particular order on the molecule, with the particular order of

the nucleotides carrying genetic information, similar to the order of letters in a word. Mendel also understood that the genetic representation of an individual differentiated between the genetic code and the outward manifestation [153].

The guiding principle of molecular biology states that information which is located in DNA is copied, through *Transcription*, to RNA. Following this, *Translation* takes place which uses RNA to assemble proteins. Transcription transforms a gene into an RNA molecule that is complementary to one strand of the DNA double helix. Following this, translation uses the information stored in three types of RNA to manufacture a protein by combining specified amino acids. RNA acts as a link between gene and protein and is complementary to one strand of the double helix, called the template strand, with the other strand being the coding strand. More recent research suggests that there are many exceptions to this, i.e. that much of the DNA does not encode proteins and instead encodes various types of functional RNAs [15].

The feature of a cell can be viewed as a *trait* and traits which are inherited are known as genotypes. Other traits emerge from the interaction between our genes and the environment in which they operate and are known as phenotypes. That is, the genotype contains the genetic code and the phenotype is the manifestation of the code coupled with the environment. Living organisms, when viewed from the perspective of evolution can be seen as a result of the mapping from their genotype to their phenotype. Evolution is the process of searching the genotypic and phenotype spaces through the use of genetic operators such as mutation and recombination followed by selection. While genetic operators operate at the genotypic level, fitness-based selection operators function at the phenotypic level. Because of this, the relationship

between the genotype and the phenotype can be expressed as the genotype-phenotype map [178].

Phenotypic variation is essential for evolution and a phenotype's variational properties are critical in an evolutionary process incorporating natural selection [192]. For adaptation to take place, changes at a genetic level have to enable *adaptive phenotypic changes* [192]. The genotype-phenotype map is an essential component in producing adaptive changes. Adaptation takes place when favourable mutations occur and this is dependant upon genetic variation which *maps onto phenotypic variation* [192]. The concepts of *Variation* and *Variability* need to be differentiated. Variation can be described as the difference between individuals in a population and can be seen as relating to a collection [151]. Variability, on the other hand, can be described as the leaning to vary and the *variability of a phenotypic trait describes the way in which it changes in response to environmental and genetic influences* [192]. A genome's variational properties, that is the variation which occurs in the phenotype as a result of genetic variation, are fundamental to evolution and the genotype-phenotype map *determines the variability of characters* [192].

The genotype-phenotype mapping can be described by pleiotropy and polygeny. Pleiotropy occurs when a single gene may simultaneously influence a number of phenotypic traits. Polygeny defines how the simultaneous interaction of a number of genes impacts on a single phenotypic trait. When pleiotropy and polygeny are present in a representation it prohibits the simplification of the genotype-phenotype mapping as a single change in the genotype may result in a number of changes in the phenotype and individual phenotypic traits are not linked to a single change at the genotypic

level [56].

1.3 Modularity & GA Representation

Modularity is a common feature of *organismic design* [192] and can be viewed as a genotype-phenotype map that contains little pleiotropic interaction between characters operating on different functions. The effects of pleiotropy are mainly seen between characters operating on a single function. A modular approach is viewed as improving evolvability by reducing the interference by *limiting the interference between the adaptation of different functions* [192]. Modularity has also been adopted in artificial systems i.e. Gruau et al. [82] and is included to simplify complexity. Overcoming the *representation problem* is critical for GAs and is fundamental as to whether or not a GA can produce adaptation. The representation problem relates to the way in which candidate solutions are represented; therefore solving an optimisation problem is only possible if the problem is encoded in a way which the variation and selection processes are effective in obtaining a solution [192]. Associated with this is adaptability and the need to consider how representations translate genotypic mutations into phenotypic ones [147].

Thus, the genetic representation of a trait determines the variability of the phenotype rather than the genetic variation within the population [192]. As the genotype-phenotype mapping is central to the representation problem, it determines the ability of the phenotype to evolve [192]. Modularity of the genotype-phenotype mapping can be described as having a separate genetic representations of character complexes, for each distinct function. [192]. Therefore, by limiting the pleiotropic effects of genes,

through the use of modularity in the genotype-phenotype mapping, the problem of unrestricted pleiotropy as the number of interactions between parts increases can be avoided [192].

Representations control the assignment of genotypes to phenotypes and represent a solution to a problem. Changing the representation can change the nature and difficulty of a problem. As the representation is the mapping from the genotype to the phenotype, for both the genotype space and the phenotype space the distances between individuals is measurable [155]. In relation to a genotype-phenotype mapping, a non-trivial map can be viewed as having the following characteristics: firstly, a phenotype can be encoded by many genotypes and secondly, the variability of a number of phenotypes will depend on their genotype [182].

Representations are viewed as being synonymously redundant when genotypes representing the same phenotype are quite alike. Also, if neighbouring genotypes correspond to neighbouring phenotypes, a representation is viewed as having a high locality of reference [154]. Non-synonymously redundant representations, on the other hand, allow genotypes which are quite different from one another represent the same phenotypes. Because of this, non-synonymous representations prevent genetic operators functioning as normal and reduce the performance of GAs [155]. Theoretical models indicate that when compared with non-redundant representations, synonymously redundant representations fail to alter the performance of selectorecombinative GAs (GAs which include selection and recombination) once all phenotypes are, on average represented by a similar number of different genotypes [155]. Only when some phenotypes are over-represented *does the performance of a GA change by the use*

of a redundant representation [155]. However when taking account of neutral theory, the accumulation of neutral mutations can lead to new paths being found [35, 37]. But with trivial neutrality, neutral traits have no impact on phenotypic evolution [183]. Non-trivial neutrality, on the other hand, suggests different genotypes within a neutral network can induce different phenotypic distributions [147], impacting on the evolutionary path.

1.4 Open Research Questions

Although there has been a lot of research undertaken into neutrality in the Evolutionary Computation community over the past number of years, the majority of this research involves GAs, possibly because of the type of representation employed by GAs. There are still many questions left unanswered, relating to variational topology and the evolutionary path, for which it is hoped that the dissertation will develop a better understanding. Open questions relating to the efficacy, variation and variability of a modular based, fixed, non-trivial many-to-one genotype-phenotype mappings and GAs are outlined below:

- How does the introduction of neutrality, through the use of a biologically inspired genotype-phenotype mapping, impact on a population's evolutionary trajectory over static and dynamic fully deceptive landscapes?
- How, over landscapes of varying degrees of hardness, does the inclusion of neutrality into the genotype-phenotype map of a GA influence genotypic and phenotypic variation?

- How does the introduction of a neutral representation into the primary structure of a GA impact on heterogeneity over a fully deceptive changing landscape?
- How does altering the granularity of neutrality, which introduces varying degrees of neutral drift, impact the phenotypic variability of a genetic algorithm with a many-to-one genotype-phenotype map?
- How does the introduction of neutrality, through the use of a biologically inspired genotype-phenotype mapping impact on the evolutionary path and phenotypic variability over a fully deceptive changing landscape?

1.5 Dissertation Goals and Hypothesis

This dissertation examines the inclusion of a multi-layered genotype-phenotype mapping in a GA based on the principles of modularity and neutral theory. The multi-layered genotype-phenotype mapping introduces redundancy into the haploid representation through an adaptation of the biological process of Transcription and Translation. The mapping allows sections of the genotype encode traits similar in concept to various types of functional RNAs, and also includes an interpretation of *Missense* mutation within the representation's layers. The Multi-layered GA (MGA) presented in the dissertation contains a modular, fixed, non-trivial representation, designed to combine Darwinism and neutral theory. Because of the nature of the mapping, it can be considered synonymous as genotypes representing phenotypes are relatively close to one another, thereby allowing variation operators to function in a problem independent manner.

By developing a GA which incorporates the principals of Darwinism and Neutral theory, the goal is to develop and analyse a biologically inspired modular, fixed, non-trivial, multi-layered genotype-phenotype mapping for a GA based on the principle of modularity and examine the impact of the representation on variation, variability and the evolutionary trajectory. To examine the impact of the representation, comparisons are performance based and a search strategy is viewed as beneficial if optimisation is achieved in fewer generations (taking into account the No Free Lunch Theorem [202]). The hypotheses emerging from the research questions are as follows:

- MGA Efficacy Hypotheses

- H_1 : *Including a biologically inspired, multi-layered, many-to-one genotype-phenotype map into a GA, benefits searching fully deceptive changing landscapes.*
- H_2 : *The inclusion of a biologically inspired, modular, fixed, non-trivial, multi-layered genotype-phenotype map into a GA, can perform as robustly as a standard GA in optimising De Jong's test suite.*

- MGA Variation Hypotheses

- H_3 : *Diversity, implicitly maintained by a many-to-one, genotype-phenotype mapping of a GA, implemented by an interpretation of the biological processes of Transcription and Translation, is beneficial in searching noisy and dynamic landscapes.*
- H_4 : *Introducing neutrality into the representation of a GA, through a modular, fixed non-trivial mapping, is beneficial in searching static landscapes*

with varying levels of difficulty.

- H_5 : *Including a modular, fixed non-trivial genotype-phenotype map which introduces a more flexible phenotypic structure and a higher degree of phenotypic variation through the use of neutrality, benefits the optimisation of solutions over dynamic landscape problems.*
- H_6 : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, maintains heterogeneity and is beneficial in promoting exploration over a fully deceptive changing landscape.*

- MGA Variability Hypotheses

- H_7 : *Altering the level of granularity alters the size of the Translation table and impacts on the search over more difficult dynamic landscapes.*
- H_8 : *The benefit of including a Transcription phase in a many-to-one, genotype-phenotype mapping, increases as the level of problem difficulty increases over more difficult dynamic landscapes.*
- H_9 : *An interpretation of Missense mutation, included within the layering of a biologically inspired, multi-layered genotype-phenotype mapping GA, assists searching more difficult dynamic landscapes.*
- H_{10} : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, alters the evolutionary trajectory and is beneficial over a fully deceptive changing landscape.*

1.5.1 Motivation

The aim of this dissertation is to develop and empirically analyse a modular, multi-layered GA, which introduces neutrality in the representation through an adaptation of the biological concepts of *Translation* and *Transcription* and uses a fixed, non-trivial genotype-phenotype map that introduces a more flexible phenotypic structure and a higher degree of phenotypic variability. This framework will allow an investigation into the impact of neutrality over various problem landscapes in order to answer the research questions posited. The objective is to investigate whether a GA incorporating a non-trivial neutral representation induces phenotypic variability and whether it can prove beneficial in searching a test suite of static and dynamic fitness landscapes, specifically examining the impact of problem difficulty on population dynamics by comparing the performance with that of a simple GA. The objective of the thesis is expressed as a series of tasks, each represented by a number of experiments.

1.5.2 Methodology

The methodology chosen is a deductive approach, as the research sets about testing the theoretical proposition laid out in the hypothesis. In other words, it moves from the general to the specific. This follows a *top-down* approach as the research moves from theory to creating testable hypotheses. To test the hypotheses a series of experiments are conducted on a test suite of suitable problems. Following this, the results produced will be quantitatively analysed to prove or disprove the hypotheses.

1.5.3 Contribution

By using a modular approach to combine Darwinism, Neutral theory and an adaptation of a number of biological concepts, the main contribution is the development of a novel, non-trivial, layered representation which introduces a tunable degree of neutrality into the primary structure of a GA in order to provide effective pressure to maintain useful diversity within the population, thereby allowing recombination cross the diversity with other building blocks, permitting continued exploration [75]. The layered representation is designed to maintain variation and increases variability, while allowing a haploid primary structure and standard variational operators function in a problem independent way. The implementation of this framework allows an examination into the impact of neutrality over specific landscapes and tests to see if the proposed layered representation impacts on the search strategy by altering genetic drift and slowing convergence, thereby locating a balance between exploration and exploitation. A secondary contribution, is in the use of the proposed representation for exploring the effects on variation and phenotypic variability of including biologically inspired variational operators within the layers of a GA, thus facilitating the examination of their impact on variation and variability. The thesis contains extracts from the following publications by the author: [86, 87, 88, 89, 90, 91, 92, 93].

1.6 Thesis Structure

The remaining dissertation layout is as follows: Chapter 2 reviews the literature with regard to the fundamentals of GAs and the measurement of GA performance

through the use of test suites. Chapter 3 examines the literature in relation to neutrality and the use of synthetic neutrality in artificial systems. Chapter 4 outlines the motivation and design of the proposed multi-layered GA (MGA). Chapter 5 examines the efficacy of the MGA and tests hypotheses H_1 and H_2 through the use of experiments designed to measure GA performance. Chapter 6 tests hypotheses H_3 , H_4 and H_5 through experiments which look at the variation or diversity (the average distance between individuals in a population [12]) associated with the MGA, Chapter 6 also introduces the *missense* mutation operator. Following this, Chapter 7 presents experiments which test hypotheses H_7 , H_8 and H_9 while looking at the arity of the MGA. Chapter 8 examines the population evolution of the MGA and looks at phenotypic distribution and population diversity, carrying out experiments to test hypotheses H_6 and H_{10} . Finally, Chapter 9 discusses the overall findings in relation to previous research and provides a summary and conclusion.

1.7 Chapter Summary

This chapter introduced and motivated the topic in question, which is to analyse a biologically inspired modular multi-layered genotype-phenotype mapping, which will allow for the examination of variation and variability. The chapter briefly outlined the areas to be discussed subsequently in the dissertation. The chapter also outlined the dissertation goals, hypothesis, motivation, methodological approach, contribution and dissertation structure. Chapter 2 gives a brief overview of the fundamentals of GAs and describes a number of well known problem landscapes which were used for experimentation in the dissertation.

Chapter 2

Genetic Algorithm

Fundamentals & Performance

Measurement

2.1 Introduction

This chapter provides an introduction to Genetic Algorithms (GAs) and an overview of the workings of GAs, discussing concepts such as schema; the schema theorem; the building block hypothesis and implicit parallelism. The chapter outlines the use of exploration and exploitation in the search process and describes the notion of a fitness landscape. The chapter then goes on to outline a number of test suites used to measure the performance of GAs. The chapter is laid out as follows; Section 2.2 briefly introduces GAs; Section 2.3 discusses GA theory and discusses GA fundamentals. Section 2.4 describes the process of searching and Section 2.5 outlines the concept of

a fitness landscape. Section 2.6 outlines a number of landscapes commonly used in the analyse of GA performance and chosen for this dissertation. Finally, Section 2.7 concludes and summarises the chapter.

2.2 Genetic Algorithms - Above the Surface

Genetic Algorithms (GAs) [95, 67], are search algorithms based on the Darwinian principal of survival of the fittest [39]. An initial population of individuals is created, each representing a possible solution to a given problem and individuals in the population are allocated a fitness score based on their suitability towards the environment in which they exist. The individuals are then subjected to environmental pressure and based on their fitness scores, selection takes place, with fitter candidates being offered a higher probability to seed the next generation through the use of operators such as recombination and mutation. In this way a new population is created and represents a new generation. The main focus in this dissertation is on simple GAs (SGA) as outlined by Vose [190], which include selection, recombination and mutation. The two processes at work here are the use of variation operators (i.e. recombination and mutation) to create diversity or variation in the population and selection, which encourages quality [174, 158].

In *Adaptation in Natural and Artificial Systems* John Holland, presented Genetic Algorithms as an abstraction of biological evolution and developed a theoretical framework, which can be adopted for GAs [95]. GAs are search techniques based on Darwin's theory of natural selection and evolution. The ability to artificially mimic, at a very basic level, the mechanics of natural selection and genetic recombination

has proved very useful in producing acceptable solutions to problems where there is very little information available with regard to the objective function, that is to say, problem domains where random search would be of little benefit in finding an adequate solution. This branch of problem solving is often referred to as *black-box optimisation*. With an encoding, which correctly represents an optimisation space, a GA does not require any other information to evolve a solution to a given problem.

GAs operate on populations of individuals, each representing a potential solution to a given problem. Each individual made up from a number of *genes* to create a *genome*. The individuals in a population are assigned *fitness* levels based on how well adapted they are to their search space, or their environment, defined by the fitness function. The idea being to introduce *selection* into the equation. Individuals then simulate reproduction, based on their level of fitness, by recombining genetic material from two individuals (*recombination*). Other natural operations such as *mutation* can also be included, normally at a much lower rate than that of reproduction. Each cycle creates a new population based on the genetic material from the old population, with new material possibly being introduced by the mutation operator. Each of these cycles is known as a *generation*. This process of selection, reproduction and mutation continues until the population reaches either a pre-described number of generations or some other external criteria is achieved.

The simple GA (SGA) [95, 67, 190] uses 1-point crossover, where two mating chromosomes are each cut once at corresponding points, and the segments following the cuts are exchanged [17]. The 1-point crossover operator randomly chooses a locus and exchanges the parts of the strings before and after that locus between

two chromosomes to create two offspring. For example, the strings *10000100* and *11111111* could be crossed over after the third locus in each to produce the two offspring *10011111* and *11100100*. Crossover is often viewed as the predominant operation in GAs. The SGA uses a single-bit mutation operator which changes the values of a location on the chromosome. It randomly flips some of the bits in a chromosome. For example, the string *00000100* might be mutated in its second position to yield *01000100*. Mutation can occur at each bit position in a string with some probability, usually very small (e.g., 0.001) [131].

Traditionally, mutation is seen as a ‘background’ operator responsible for re-introducing mistakenly lost gene values (alleles), preventing genetic drift, and providing a small element of random search when the population has largely converged. It is generally accepted that crossover is the main force leading to a methodical search of the problem space. However, examples in nature show that asexual reproduction can evolve sophisticated creatures without crossover [17]. Tate and Smith [181] argue that the optimal mutation rate depends strongly on the choice of encoding, and that problems requiring non-binary encoding may benefit from mutation rates much higher than those generally used with binary encodings. While Bäck argues that variation in the mutation rate may accelerate optimisation [6, 7].

De Jong’s work *An Analysis of the Behaviour of a Class of Genetic Adaptive Systems* [43] set the stage for many of the heuristics which have been used in the creation of GA theory. De Jong developed a basic structure and a number of operations which created the SGA, consisting of roulette wheel selection (fitness-proportionate), one-point crossover and mutation. These three processes worked on a population of

binary strings with fixed orderings and with each generation a certain percentage of the population, known as the generation gap, would be replaced by the processes. De Jong's work defined GAs as a group of algorithms parameterised by the population size (n), the probability of crossover (p_c), the probability of mutation (p_m), and the generation gap (G).

Representation issues for GAs tend to address the question of how to engineer GAs. Related to representation issues is the choice of genetic operators for introducing variation into a population. One reason that binary linearly ordered representations are frequently used is that standard mutation and crossover operators can be applied in a problem independent way. "*In natural genetic systems, one gene (approximately) codes for the one protein regardless of where it is located, although the expression of a gene (when the protein is synthesised) is indirectly controlled by its location*" [132]. Despite this, most current GA implementations use a simple binary alphabet, linearly ordered along a haploid string. Researchers interested in engineering applications have believed that the use of "high-cardinality alphabets" including real numbers as alleles should be advocated [132].

2.3 Genetic Algorithms - Below the Surface

In order to understand how GAs work, an understanding of Holland's notion of *building blocks* [95] is essential [69]. Holland puts forward the *schema theorem* [95] to explain how a GA can engage in a complex and robust search by implicitly sampling hyperplane partitions of a search space [199]. A schema is a collection of gene values which may be represented (in a binary coding) by a string of characters contained

within the alphabet 0, 1, *, where * means “don’t care” [133]. A chromosome contains a particular schema if it matches that schemata, with the “*” symbol matching anything. For example, the chromosome “1010” contains many schemata including, “10**”, “*0*0”, “**1*” and “101*”. The order of a schema is the number of non-* terms it contains (2,2,1,3 respectively in the above examples). The defining length of a schema is calculated by the distance between the outermost non-* symbol (i.e. 2,3,1,3 respectively from the examples above).

Holland described mathematically how the number of representatives of a schema alter as one generation progresses to the next. This was summarised by Holland in the *Schema Theorem* [95].

$$m(H, t + 1) \geq m(H, t) \frac{\bar{f}(H, t)}{\bar{f}(t)} \left(1 - p_c \frac{\delta(H)}{\ell - 1} - p_m o(H)\right) \quad (2.1)$$

Where $m(H, t)$ is the number of occurrences of schema H at time t . $\bar{f}(H, t)$ represents the average fitness of the occurrences of schema H at time t . $\bar{f}(t)$ represents the average fitness of the population at time t . $\delta(H)$ is the defining length of schema H , p_c is the probability of crossover and p_m is the probability of mutation. ℓ represents the length of the string and finally $o(H)$ is the order of schema H [67].

The schema theorem helps to explain the power of a GA in terms of how schemata are processed. Members of the population are individually given the chance to reproduce, and create offspring. The number of chances an individual receives is directly related to that individual’s fitness; in other words, there is a high probability that fitter individuals will contribute more of their genes to the next generation. It should be noted that it is assumed that a individual’s high fitness is due to the fact that

it contains good schemata. By passing ‘fitter’ schemata to the next generation it is more likely that a better solution will be found. Others such as Stephens and Waelbroeck [179, 180], argue that through course graining, there is no preference for short, low-ordered schemata and there can exist a preference for large schemata.

Holland [95] showed that the optimal way to explore the search-space is to allocate reproductive opportunities to individuals in proportion to their fitness relative to the other creatures in the population. By doing this, good schemata receive an exponentially increasing number of chances to reproduce in future generations. Holland also showed that as each individual contains many different schemata, the number of schemata effectively being processed in each generation is of the order n^3 (n being the population size), this property is called *implicit parallelism* [16].

The traditional design focus for crossover operators has been combination, random parts (schemata) from both parents are crossed into an offspring solution. This design basis is a legacy of the original schema theorem. The focus is on implicit parallelism and the idea of sampling of schemata, particularly low-order schemata (building blocks) [81], with crossover viewed as the predominant operator. However, it should also be noted that crossover and mutation do not operate independently of one another and interact with the other elements of an evolutionary system [46]. In relation to mutation many biologists see it as the main source of raw material for evolutionary changes [111]. This theory moves away from the traditional view that low mutation rates are used in GAs as they tend to lead to an efficient search of the solution and that high mutation rates result in diffusion of search effort and premature extinction of favourable schemata in a population. Mutation is primarily

useful for retrieving valuable schemata, this implies the belief that the mutation rate ought to be set so that the rate at which highly fit schemata are accidentally deleted from the population is roughly equal to the rate at which mutations introduce new desirable schemata. Implicit in this view is that it is assumed there is no need to find highly fit schemata that are not present in the initial population, or that cannot be generated from that population by repeated breeding operations. Put another way, this view assumes that nearly all the useful gene alleles are present in the initial population.

While the Schema theorem indicates that selection emphasises fit schemata and some of these schemata are destroyed by genetic operators, the overall lesson is that low order schemata of short defining length, that is, highly fit schemata which are not too disrupted by genetic operators tend to increase from generation to generation. The schema theorem can also be written as

$$m(H, t + 1) \geq m(H, t)\phi(H, t)[1 - \epsilon(H, t)] \quad (2.2)$$

which is independent of the particular choice of GA operators. Interpretation is also made easier as the effect of selection is isolated by the reproduction ratio $\phi(H, t)$ and the effect of the genetic operators is given by the distribution factor $\epsilon(H, t)$. A schema can grow or decay according to the net growth factor $\phi(H, t)[1 - \epsilon(H, t)]$ [38].

The *Building Block Hypothesis* argues that the power of a GA lies in being able to find good *building blocks* [67]. These building blocks are schemata of short defining length consisting of bits that work well together and tend to improve performance when incorporated into an individual. A successful coding scheme encourages building blocks to form by ensuring that, firstly, related genes are close together on the

chromosome and secondly, there is little interaction between genes. *Interaction* (also called *epistasis*) between genes means that the contribution of a gene to the fitness depends on the fitness of other genes in the chromosome. Epistasis refers to any kind of strong interaction among genes. There is always some epistasis between genes in multimodal fitness functions, which are important in GA research as unimodal functions can be solved using simpler methods. If both related genes are close together on the chromosome and little interaction between genes is observed, then GAs will be as efficient as predicted by the schema theorem. However, genes may be related to one another in ways that do not allow all the closely related genes to be located in close proximity in the string. Also, if the designer does not know exactly the relationship between the genes, this may prevent the creation of a successful coding scheme. If there is little interaction between the genes on the chromosome then this implies that related genes are close together on the chromosome. Suppose the contribution to overall fitness of each gene was not dependent on any other genes, then it may be possible to solve the problem by hillclimbing on each gene in turn. It is generally accepted that when designing the code schemes, one should attempt to conform with Goldberg's building block hypothesis so as to ensure that the GA will perform effectively [69].

2.4 Exploring the Search Space

To illustrate the mechanism of a GA in searching a fitness landscape let us see how a GA can sample hyperplane partitions. Consider a 3-dimensional space, if a problem is encoded with 3 bits, this can be represented as a cube with string 000 at

the origin. The cube's corners are numbered bit strings and all adjacent corners are labelled by bit strings that differ by exactly 1 bit. The front plane of the cube contains all points which begin with 0. A "*" can be used as a "don't care" or wild card match symbol, by doing this the plane can be represented by the string 0** [133]. Strings which contain "*" are called *schemata*, each schemata corresponds to a hyperplane in the search space. The hyperplane refers to the number of actual bit values that appear in the schema. For example, 1** is the order-1 while 1**1*****0** would be of order-3. In other words the order of a schema is the number of non "*" symbols it contains. The defining length of a schema is the distance between the outermost non "*" symbols, i.e. 10** has a defining length of 2, *0*0 has a defining length of 3, while **1* has a defining length of 1.

A 4-dimensional space can be represented by a cube suspended inside another cube. The points in the inner cube (the 3-dimensional cube) and the outer 4-dimensional cube are labelled the same. Next, prefix the inner cube label with a 1 bit and each outer cube with a 0 bit, thus allowing the creation of adjacency in strings which are different by 1 bit. The outer cube now corresponds to the hyperplane 0*** while the inner cube corresponds to the hyperplane 1*** [95]. Thus *0** identifies the subset of points that corresponds to the front of both cubes, and the order-2 hyperplane 10** corresponds to the front of the inner cube [199].

A bit string corresponds to a particular schemata if it can be made *from the schemata by replacing the "*" with the appropriate bit value* [199]. Bit strings matching a certain schemata are part of the hyperplane partition represented by that specific schemata [199]. *"Every binary encoding is a chromosome which corresponds to a cor-*

ner in the hyperplane and is a member of $2^L - 1$ different hyperplanes, where L is the length of the binary encoding” [199]. The string of all “*” symbols relates to the space itself and is not counted as a partition of the space [95]. Each string has L positions in it and each position can be either a bit value contained in the string or the “*” symbol [199].

Optimisation algorithm often uses two mechanisms to find a global optimum; *exploration*, which investigates *new and unknown areas in the search space* [16], and *exploitation* which makes use of information gained from locations previously visited, in order to locate better locations [16]. These requirements of exploration and exploitation are contradictory, therefore a good search algorithm must find a compromise between the two [16]. A random search can be used for exploration but does no exploitation. On the other hand a ‘*hillclimbing*’ method excels at exploitation, but carries out very little exploration. Combining exploitation and exploration can be effective but a balance needs to be struck between the two, which can be difficult to achieve, i.e. what amount of exploitation is performed before giving up and exploring further [16]. A GA combines both exploitation and exploration at the same time [95].

Although this is true in theory, there are a number of problems which arise in practice. These problems occur due to certain assumptions made for simplification purposes. Some of these assumptions include: an infinite population size; the fitness function accurately reflecting the problem and also that there is little interaction between genes in a chromosome [16]. The infinite population size can *never be satisfied in practice, because of this the performance of a GA will be subject to stochastic errors* [16], i.e. *genetic drift* (which is also present in nature).

This implies that even if there were no selection pressure, in other words a constant fitness function, members of a population will still manage to converge to some point in the solution space due to the accumulation of stochastic errors [16]. If a gene becomes predominant in a population, there is every likelihood that it will become more predominant in the next generation. There is also a chance that it will become less predominant in the next generation. But if the increase in predominance is sustained over a number of successive generations, and the population is finite, then a gene can spread to all creatures in a population. Once a gene has converged in this manner, it is fixed i.e. crossover cannot introduce new gene values [16]. The knock on effect is that genes eventually becomes fixed. Thus *the rate of genetic drift provides a lower-bound on the rate at which a GA can converge towards the correct solution* [16]. In other words if a GA is to exploit gradient information in the fitness function, the fitness function must provide a slope sufficiently large to counteract [16] any trend towards genetic drift. The rate of genetic drift can be reduced by increasing the mutation rate. However, if the mutation rate is too high, the search becomes effectively random; therefore the gradient information in the fitness function is not exploited [16].

2.5 The Fitness Landscape

Shortly after the first mathematical models of Darwinian evolution were developed, Sewall Wright [203] recognised a property of population genetic dynamics, that is, when fitness values of genes are evaluated, the genetic makeup of a population can evolve into multiple domains of attraction. The specific fitness interaction is

known as *epistasis*, where the effect on fitness from altering one gene depends on the allelic state of other genes. Epistasis makes it possible for the population to evolve towards different combinations of alleles, depending on its initial genetic combination. Thus, Wright discovered a conceptual link between a microscopic property of organisms (fitness interactions between genes) and a macroscopic property of evolutionary dynamics, multiple population attractors in the space of genotypes [2]. *Implicit in this idea is a collection of genotypes arranged in an abstract metric space, with each genotype next to those other genotypes which can be reached by a single mutation, as well as a value assigned to each genotype* [177].

Wright illustrated this by using the metaphor of a landscape of multiple peaks, in which a population would evolve by moving up hill until it reached its local fitness peak. This visualisation of the *adaptive landscape* is the term used to describe multiple domains of attraction in evolutionary dynamics. Wright looked specifically at how populations could get away from local fitness peaks to higher ones through stochastic fluctuations in small population subdivisions. This was one of the first conceptions of stochastic processes for the optimisation of multimodal functions [2]. Multimodality or the number of peaks in a search space, is an important characteristic of that particular search space [110].

The idea of a fitness landscape is used as a framework for reasoning about evolution. Biological organisms can be characterised by their genotype, which is the genetic ‘encoding’ of the organism, or their phenotype, which is the actual form and behaviour of the organism. A fitness value can be assigned to each phenotype, which denotes its ability to survive and reproduce [186]. Evolution can be viewed as a

process that searches, by means of genetic operators like crossover and mutation, a fitness landscape of possible genotypes, looking for genotypes that encode highly fit phenotypes [44]. Evolution searches for solutions, encoded in genotypes, to find fit organisms which are capable of reproduction. Every genotype will have a relative fitness assigned to it. This is determined by a fitness function. The fitness landscape is then constructed by assigning the fitness values of the genotypes to the corresponding points in the genotype space. To visualise this picture, each point in the genotype space is given a ‘height’ according to its fitness. From this a ‘mountainous’ landscape is formed, where the highest peaks designate the best solutions. A local optimum, or peak, in such a landscape is defined as a point that has a higher fitness than all its neighbours [97]. *In evolutionary computation the notion of evolvability refers to the efficiency of evolutionary search. It has been shown that the structure of a fitness landscape affects the ability of evolutionary algorithms to search. Three characteristics specify the structure of landscapes. These are the landscape smoothness, ruggedness and neutrality. The interplay of these characteristics plays a vital role in evolutionary search* [104]. Much research has been carried out examining the relationship between the landscape structure of problems with known difficult and GA performance, including Grefenstette [81], Lipsich [120], Manderick [127], Goldberg [66] and Horn [98], facilitating the use of well understood fitness landscapes to be used in the empirical analyses of GAs.

To describe fitness landscapes, a notion of a distance between genotypes is needed. Genotypes are codings, and different codings can cause different distance measures. Furthermore, there is often more than one distance measure, or metric, which can be

defined for one and the same coding. Usually, a coding in the form of bit strings is used. The Hamming distance [83] between two bit strings is defined as the number of corresponding positions in these bit strings where the bits have a different value. So, the distance between 010 and 100 is two, because the first and second positions have different values. A normalised Hamming distance can be defined by dividing the Hamming distance between two bit strings by the length of these bit strings. By adopting this approach, the distance measure is independent of the length of the bit strings. A normalised Hamming distance of 0.5, for example, means that half the bits of two bit string have a different value [96].

2.6 Measuring GA Performance

In order to evaluate the changes in variation and variability associated with the proposed layered genotype-phenotype map, a variety of problem landscapes were used in the dissertation. The problems chosen consist of both static and dynamic landscapes. The static problems include variations of the OneMax problem [159], various deceptive landscapes [198, 73] and De Jong's test suite [43]. The dynamic landscapes used include a number of changing deceptive landscapes [87] and a changing Sphere model landscape [94]. A description of the various problem landscapes used is given in this section.

2.6.1 Static Landscapes

Deceptive Landscapes

Bethke [20], examined the identification of problem difficulty within GAs. Bethke's approach revolved around converting the problem in question into a domain where a range of small-order schemata averages were easier to establish. With this approach predicting whether or not small highly fit schemata would eventually find the global optimum was made easier. Bethke converted problems into the domain of Walsh function coefficients. By adopting this approach an order- k schema's average fitness is easier to conclude as originally a large number of individuals in a population would have to be examined, whereas reviewing only 2^k Walsh coefficients will produce same. Bethke hypothesised that if the scale of the Walsh coefficients decreased quickly, in association with increasing order, then small schemata would determine the search space and the problem would be easy for the GA. Bethke, by reversing his transformation, could design problems which were difficult for a GA, as he hypothesised that the higher the minimum order of a schemata required to locate the global optimum, the more difficult the problem. Bethke also recognised misleadingness and isolation and pointed out that with GA-hard functions "*not only must a GA-hard function be epistatic, but the epistasis must be misleading*" [20]. In other words short, high-performance schemata must point toward poor areas of the space. In many ways Bethke's work paved the way for the development of test functions which allowed for the evaluation of GAs.

With regard to testing the performance of GAs, a number of authors have used a class of problem considered *GA-Hard* [70]. A problem is considered hard in many

cases if it is difficult for the GA to find the optimum; in other words we are looking for a single optimum point surrounded by points which have a relatively low fitness level. Problems are often created by taking advantage of the GA so as to allow selection to deliberately lead the search away from the optimum. By using a binary encoding of a solution, GAs can search complex landscapes by sampling hyperplanes in an n -dimensional hypercube. GAs alter the rates of sampling of hyperplanes in relation to encodings of the solution space. In relation to deceptive problems hyperplanes drive the search away from the global optimum and towards a local optimum [70].

Deceptive functions can be viewed as having two optima: a global optima, located at position $1, 1, \dots, 1$ and a local optima, located at position $0, 0, \dots, 0$. If the fitness value of the global optimum is $f_{global} = x$ and the fitness value of the local optimum is $f_{local} = x - 1$, then the remaining points in the landscape are given fitness values based on their distance from one of the optima in a way that the basin of attraction for the global optimum is significantly smaller than the basin of attraction of the local optimum. In other words the search is designed to draw you away from the global optimum. It is this feature which makes the problem deceptive. Many authors have used deceptive landscapes as a test bed for GAs including Goldberg [68, 72, 50], Kargupta [105], Grefenstette [80], Jones [102, 103] and Dasgupta [40] to name but a few.

The Building Block hypothesis is one of the fundamental theories of GAs. By combining highly fit low-order schemata into higher-order schemata a GAs seek near-optimum performance. However, highly fit low-order schemata may not combine to form highly fit high-order schemata. Goldberg [65], by devising a 2-bit problem where

the combination of the best order-1 schemata did not produce the global optimum, introduced the notion of deception to mislead a simple GA (SGA). Goldberg constructed this *minimal deception problem* (MDP) by analysing the schema averages directly, as opposed to using Walsh coefficients similar to Bethke, and analysed how the GA handled the problem. Goldberg's results indicated that the GA could overcome this level of problem difficulty, but occasionally a loose ordering of this problem, or large defining lengths, could combine with a poor starting population to frustrate the GA. Deceptive problems cause difficulties for GAs because firstly, the global solution is isolated and secondly, information deceives the GA to locate sub-optimal solutions. By applying the methodology developed by Goldberg to higher-order building blocks, a class of problems described as deceptive problems has emerged. Problems of bounded difficulty, which have become effectively the test bed for GAs, resulted from the ability of deceptive problems to be set within larger problems.

3-bit Deceptive Problem - Test 1 & Test 3

A fully deceptive problem of order-N can be viewed as being deceptive when all of the lower-order hyperplanes lead away from the global optimum and towards a deceptive attractor [198]. To illustrate deception, let us consider a fully deceptive order-3 function, where the information about the hyperplane which the order-1 and order-2 schemata represent in the search space direct the search away from the global optimum towards a deceptive attractor [198]. If the bits 111 represent the global optimum and the bits 000 represent the *deceptive attractor*, then a full order-3 deception would be similar to that defined by Goldberg, Korb and Deb [73] which is illustrated

Order-3 Problem			
String	Fitness	String	Fitness
f(000)	28	f(010)	22
f(110)	0	f(101)	0
f(001)	26	f(100)	14
f(011)	0	f(111)	30

Table 2.1: Deceptive Order-3 Problem - Fitness Values

in Table 2.1 and shows the fitness values for each bit string.

10 3-bit Fully Deceptive Problems - Test 1, Test 3 & Test 5

One failing of the 3-bit fully deceptive problem is that it is too small to really demonstrate a search strategy. The 10 3-bit problem as outlined by Goldberg [73] expands the 3-bit problem into 10 3-bit deceptive order-three sub-functions. The effect of this is to make it difficult for the SGA to solve as it tends to converge prematurely, with the sub-functions being drawn towards the deceptive attractor rather than the global optimum [73]. However, in Goldberg's 10 3-bit function each 3-bit sub-function is associated with the adjacent 3-bit sub-function. In this way the first sub-function is associated with bits one to three and the second sub-function is associated with bits four to six and so on. This function is relatively difficult using most functional optimisation standards and the search space is quite large.

To increase the level of difficulty, a *loose ordering* has been included, which makes the problem fully deceptive. This is achieved by increasing the defining length to twenty, where the defining length is the maximum distance between two defining symbols in a schema. The effect of this is to make it difficult for the simple GA to solve as it tends to converge prematurely, with the sub-functions converging on 000

rather than the optimum of 111 [73]. For example, rather than having bits one to three linked together to create a sub-function, increasing the defining length to twenty means our sub-function now consists of bits one, eleven and twenty one for the first sub-function and bits two, twelve and twenty two for the second sub-function etc. A count of the number of sub functions discovered throughout the search is carried out, with ten being the maximum number achievable. By increasing the 3-bit problem into ten 3-bit sub-functions, as outlined, we create a 10 3-bit fully deceptive landscape which through the increase in defining length for each sub-function, dramatically increases the level of problem difficulty.

30 3-bit Fully Deceptive Problems - Test 5

Similarly by increasing the ten 3-bit problem into 10 3-bit sub-functions, as outlined, we create a 30 3-bit fully deceptive landscape which through the increase in defining length for each sub-function, thereby dramatically increasing the level of problem difficulty. The 30 3-bit fully deceptive problem is an extension to the 10 3-bit fully deceptive problem designed to increase the problem difficulty as the search space is increased, this approach to increasing difficulty was also adopted by Pelikan [146].

30-bit & 90-bit OneMax Problem - Test 5

The OneMax Problem [159] is a simple problem consisting of maximizing the number of ones in a bit-string. Formally, this problem can be described as finding a string $\vec{x} = \{x_1, x_2, \dots, x_N\}$, with $x_i \in \{0, 1\}$, that maximizes the following equation:

$$F(\vec{x}) = \sum_{i=1}^N x_i \quad (2.3)$$

As the OneMax bit counting problem is a relatively easy problem landscape two variations were used in Test 5, a 30-bit OneMax problem and a more difficult 90-bit OneMax problem.

De Jong's Testsuite - Test 2

De Jong [43] created test functions and performance measures to determine the effect various modifications to the parameters would have on the GA. Five test functions were designed to access the ability of a simple GA (SGA) over various landscapes. The functions were chosen by De Jong because they represented many common difficulties found in optimisation problems. De Jong's test suite was originally designed as a minimising problem and can easily be changed into maximising problems. The test functions accessed SGAs over problems which were continuous or non-continuous, convex or non-convex, unimodal or multimodal, quadratic or non-quadratic, deterministic or stochastic, and finally low-dimensional or high-dimensional. With this test bed created, De Jong's experiments involved running a SGA and altering its parameters. The results distinguished between what De Jong termed *on-line* and *off-line* optimisation. On-line referred to a problem which required a continuing level of optimisation, whereas off-line referred to a problem where the solution discovered after a predetermined amount of time could be used.

By running a SGA and altering its parameters over the test bed, De Jong's results exhibited a number of conclusions. Firstly, off-line results were more successful with

large populations while on-line performances were more successful with smaller populations. The results also indicated that higher mutation rates, $P_m > 0.1$ caused the SGA to behave in a manner similar to a random search; a probability of crossover, $P_c = 0.6$ proved a good balance between on-line and off-line performances and overlapping populations resulted in poorer performances. While investigating variants of the SGA, De Jong found that schemas with less noise, which implemented fitness-proportionate selection more stringently, outperformed roulette wheel selection. With regard to experiments relating to recombination operators, De Jong found that multiple point crossover operators which caused more disruption proved detrimental to the SGA. De Jong's experiments, methodologies and conclusions provided a framework for many researchers to examine GAs [78, 160, 157, 123] etc.

The Sphere Function (De Jong's Test Suite f_1)

The first function (f_1) is relatively easy to optimise as it is smooth, convex and unimodal. The minimum is located at the centre of the N-dimensional sphere and there are no local minima. This function is normally used to measure the efficiency of a particular algorithm. A graphical representation of the generalised Sphere function is shown in Figure 2.1 and the function has the following definition:

$$f_1 = \sum_{i=1}^2 x_i^2 \tag{2.4}$$

where

$$-5.12 \leq x_i \leq 5.12$$

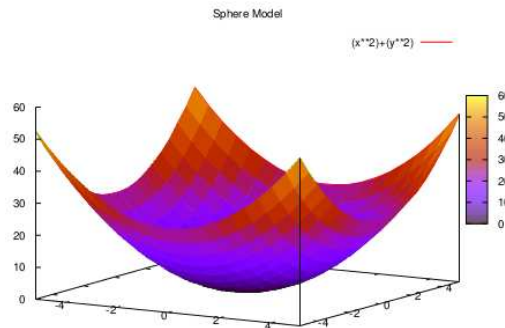


Figure 2.1: The Sphere Model.

Rosenbrock's Function (De Jong's Test Suite f_2)

The second function (f_2), Rosenbrock's function, is a frequently used optimisation problem [166] and is also known as Rosenbrock's valley or Rosenbrock's banana function.

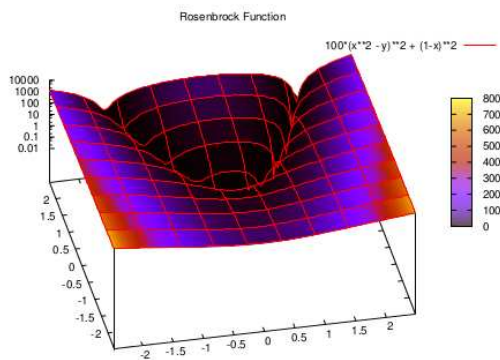


Figure 2.2: Rosenbrock's Function.

It is a two-dimensional function containing a deep valley, shaped in a parabola. Rosenbrock's function, also known as Rosenbrock's Saddle function is viewed as being a relatively difficult minimisation problem, as the algorithm is dealing with repeatedly

changing direction of the search and the difficulty lies in its ability to converge. Figure 2.2, illustrates Rosenbrock's function and the function has the following definition:

$$f_2 = 100(x_1^2 - x_2)^2 + (1 - x_1)^2 \quad (2.5)$$

where

$$-2.048 \leq x_i \leq 2.048$$

The Step Function (De Jong's Test Suite f_3)

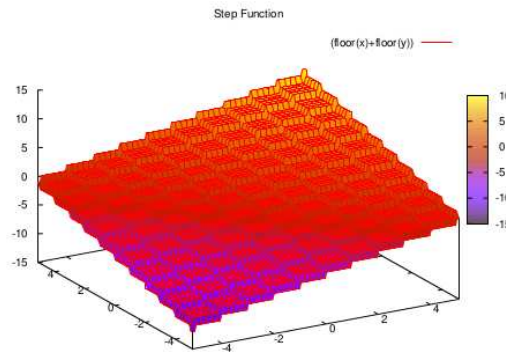


Figure 2.3: Step Function.

The third function (f_3), the Step function, exhibits many plateaus and represents problems with flat surfaces. This class of problem proves difficult for many algorithms as they don't provide information as to the most advantageous direction. Flat surfaces are obstacles for optimisation algorithms, as it is difficult to decide which direction to go without having any information available. Figure 2.3 shows a plot for the Step function. The Step function has the following definition:

$$f_3 = \sum_{i=1}^5 [x_i] \quad (2.6)$$

where

$$-5.12 \leq x_i \leq 5.12$$

Quadratic Function with Noise (De Jong's Test Suite f_4)

The fourth function (f_4) is a quadratic function which includes Gaussian noise. This function is designed to test the behaviour of an algorithm in the presence of noise. The function's global optimum depends on the expectation value of a random variable, and also includes several randomly distributed local minima. In this dissertation the experiments use a 30-dimensional function which contains noise to ensure that points return a different value each time they are evaluated. The plot for the quadratic function without noise is illustrated in Figure 2.4. The quadratic function with noise is defined as follows:

$$f_4 = \sum_{i=1}^{30} (ix_i^4 + Gauss(0, 1)) \quad (2.7)$$

where

$$-1.28 \leq x_i \leq 1.28$$

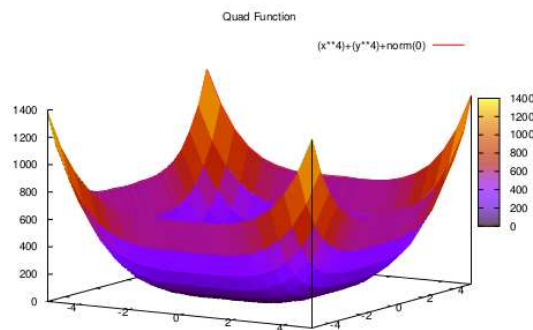


Figure 2.4: Quadratic Function.

Shekel's Foxhole Function (De Jong's Test Suite f_5)

The fifth and final function (f_5), Shekel's Foxhole, is a multidimensional, multimodal, continuous, deterministic function which contains many local optima. The difficulty lies in the algorithms ability not to get trapped in a local optima. This 2-dimensional function contains 25 different *foxholes*, each varying in depth, surrounded by relatively flat surfaces. Shekel's Foxhole function has the following definition:

$$f_5(x_i) = 0.002 + \sum_{j=1}^{25} \left(\frac{1}{j} + \sum_{i=1}^2 (x_i - a_{ij})^6 \right) \quad (2.8)$$

where

$$-65.536 \leq x_i \leq 65.536$$




Figure 2.5: Shekel's Foxholes.

2.6.2 Dynamic Landscapes

In relation to classifying dynamic landscapes, De Jong [45] proposed four different categories, each relating to real world problems. The first category involves drifting

landscapes, where the topology changes little by little over time. The second category involves landscapes that undergo significant morphological changes with peaks of high fitness become areas of low fitness and areas of the landscape which were uninteresting become new areas of high fitness. The third category proposed by De Jong involves landscapes with cyclic patterns and the fourth and final category are landscapes where changes are *abrupt and discontinuous* [45]. Branke [23] used the following criteria to characterise changing environments: frequency of change; severity of change; predictability of change and cycle length or cycle Accuracy. In this dissertation the changing landscape can be viewed as a combination of De Jong's fourth category with the change being abrupt and discontinuous and Brank's severity of change characterised by the distance between the old and new optimum.

Over the past number of years there has been a significant amount of research into the performance of EAs in tracking a moving optimum in dynamic environments including Cobb [31], Grefenstette [32, 79], Dasgupta [41], Bendtsen [19], Branke [22, 24, 23, 25], Kubalik [115]. Techniques emerging from past research can broadly be classified into two categories, with the first examining the population in relation to variation and exploration and the second which uses memory to exploit past information [85]. The focus of this dissertation in relation to dynamic environments lies with the first category, examining the balance between convergence and diversity.

Changing Landscape Problems

Natural selection requires diversity to be present in a population for adaptation to take place [122]. Diversity is critical for GAs particularly when the landscape is evolving as recombining a homogeneous population will not enable the GA to locate the new optimum [135]. To examine the performance of any GA over a changing landscape, once the GA reaches a predefined point during the search, the fitness function is altered, which has the effect of changing the landscape and creating a new global optimum [31]. By doing this, the focus is placed on the adaptive qualities of the GAs in response to altering the target of the search and outlines the algorithm's ability to maintain a balance between exploration and exploitation. Table 2.2, outlines the changes made to the fitness function shown and shows the new *deceptive attractor* 111, with 000 being the new global optimum. Again, the low level building blocks lead the search away from the optimum, in other words all of the low-order hyperplanes direct the search to the deceptive attractor [198].

10 3-bit Changing Fully Deceptive Problem - Tests 1, 3 & 5

The 10 3-bit changing fully deceptive, loosely ordered problem, described in 2.6.1, incorporates the 3-bit deceptive problem as described in [73]. However the fitness function changes at a predefined point in the search, altering the landscape. The changes to the fitness are outlined in Table 2.2, with Table 2.2a outlining the fitness values before the landscape change and Table 2.2b illustrating the fitness values after the change.

Deceptive Order-3			
String	Fitness	String	Fitness
f(000)	28	f(010)	22
f(110)	0	f(101)	0
f(001)	26	f(100)	14
f(011)	0	f(111)	30

Reversed Order-3			
String	Fitness	String	Fitness
f(000)	30	f(010)	0
f(110)	26	f(101)	22
f(001)	0	f(100)	0
f(011)	14	f(111)	28

(a) Deceptive Order-3 Problem

(b) Reversed Deceptive Order-3 Problem

Table 2.2: Order-3 Deceptive Problem & Reversed Order-3 Problem

30 3-bit Changing Fully Deceptive Problem - Test 5

The 30 3-bit changing fully deceptive, loosely ordered problem, outlined in 2.6.1, is an extension of the 10 3-bit fully deceptive problem. The increase in the solution space is designed to make the landscape more difficult to search. The 3-bit fitness values associated with this problem are those shown in Table 2.2.

Changing Sphere Model Problem - Test 4

The changing Sphere model [94], is the same as the Sphere model (f_1) discussed previously (2.6.1) with the addition that the fitness function is altered during the run. The change is designed to test the GAs ability to adapt and the details of the Sphere Model Changing Environment are as follows:

$$f_2(x(\vec{t})) = \begin{cases} \sum_{i=1}^n x_i^2(t) & : t \bmod a \text{ even} \\ \sum_{i=1}^n (x_i - b)^2 & : t \bmod a \text{ odd} \end{cases} \quad (2.9)$$


$$-5.12 \leq x_i \leq 5.11$$

$$a = 1500 \text{ generations} \quad ; \quad b = 4$$

$$\min(f_2) = \left\{ \begin{array}{l} f_2(0, \dots, 0) \quad : \quad t \bmod a \text{ even} \\ f_2(b, \dots, b) \quad : \quad t \bmod a \text{ odd} \end{array} \right\} = 0$$

4-bit Deceptive Changing Landscape - Test 6 & Test 7

In this dissertation a number of experiments are also conducted over fully deceptive Order-4 landscapes [198], which when reversed, allows the local optimum to become the global optimum. Figure 2.6a graphically illustrates the landscape of the 4-bit deceptive problem [198], with the x and y co-ordinates indicating the location on the grid. To analyse the adaptive qualities of both GAs, the landscape reverses at a predefined number of generations, with Figure 2.6b illustrating the landscape after the change and shows the global optimum changing to become the local optimum and visa-versa.



(a) 4-bit Deceptive Landscape

(b) 4-bit Reversed Deceptive Landscape

Figure 2.6: 4-bit Changing Deceptive Landscape

Table 2.3 shows the fitness values associated with each 4-bit string, with Table 2.3a

Deceptive Order-4			
String	Fitness	String	Fitness
$f(0000)$	28	$f(0001)$	26
$f(0010)$	24	$f(0011)$	18
$f(0100)$	22	$f(0101)$	16
$f(0110)$	14	$f(0111)$	0
$f(1000)$	20	$f(1001)$	12
$f(1010)$	10	$f(1011)$	2
$f(1100)$	8	$f(1101)$	4
$f(1110)$	6	$f(1111)$	30

Reversed Deceptive Order-4			
String	Fitness	String	Fitness
$f(0000)$	30	$f(0001)$	6
$f(0010)$	4	$f(0011)$	8
$f(0100)$	2	$f(0101)$	10
$f(0110)$	12	$f(0111)$	20
$f(1000)$	0	$f(1001)$	14
$f(1010)$	16	$f(1011)$	22
$f(1100)$	18	$f(1101)$	24
$f(1110)$	26	$f(1111)$	28

(a) Deceptive Order-4 Problem

(b) Reversed Deceptive Order-4 Problem

Table 2.3: Deceptive & Reversed Deceptive Order-4 Problem - Fitness Values

outlining the fitness values before the landscape change and Table 2.3b outlining the fitness after the landscape changes.

2.7 Chapter Summary

This chapter gave a brief overview of GAs and examined the workings behind GAs, describing Schema, the Schema theorem, traditional variation operators, implicit parallelism, the building block hypothesis, epistasis and genetic drift. Following this, the chapter discusses the search space and the idea of a fitness landscape, outlining a number of well known static and dynamic problems describing their use in evaluating GA performance. The problem landscapes outlined above include static deceptive landscapes, changing deceptive landscapes, De Jong’s test suite and a changing Sphere Model landscape. These landscapes were chosen as they are well understood and are commonly used to evaluate GAs. By using the landscapes described the proposed neutral genotype-phenotype map contained in the MGA was evaluated on a

variety of problems, including GA-Hard problems. The problems chosen contained a mix of characteristics, including: continuous or non-continuous problems; convex or non-convex problems; unimodal or multi-modal problems; quadratic or non-quadratic problems; deterministic or stochastic problems and finally low-dimensional problems or high-dimensional problems. The next chapter moves on to examine the concept of Neutral theory and neutrality, describing neutrality as found in nature and various ways in which synthetic neutrality is implemented in artificial systems.

Chapter 3

Neutrality

3.1 Introduction

Sewall Wright noticed that many random changes in the frequency of alleles occurring in a population were not related to selection [203]. His observations indicated that this *genetic drift* was an important component in the evolutionary process. *Neutral theory* as proposed by Kimura [111] which offered an alternative to the Darwinian view, states that the mutations involved in the evolutionary process are neither advantageous nor disadvantageous to the survival of an individual, and that most mutations are caused not by selection, but rather by random genetic drift. However Kimura [112] pointed out that although natural selection does play a role in adaptive evolution, only a tiny fraction of DNA changes are adaptive. The vast majority of mutations are phenotypically silent.

By adopting the principal of Darwinism, simple genetic algorithms (SGAs), can be viewed as implementing the process of evolution without containing any explicit

neutral mutations. In other words, each mutation is either an advantage or a disadvantage to the individual in terms of fitness, with selection then propagating the fitter individuals. As the search progresses, exploration-exploitation ratios decrease as the population converges. If we are to implement a GA based on the principles of neutral theory then, neutrality needs to be introduced. Neutrality can be viewed as a situation where a number of different genotypes can represent the same phenotype.

This chapter reviews the literature associated with neutral theory, neutrality in biology and synthetic neutrality as implemented in artificial systems. The chapter is organised as follows: Section 3.2 examines neutral theory and neutrality, discussing the development of neutral theory from a biological perspective. Section 3.3 looks at the use of synthetic neutrality and examines how the concept of neutrality has been used in artificial systems and GAs, beginning with neutrality through the use of various fitness landscapes and then neutrality implemented through genotype-phenotype mappings. Finally, Section 3.4 concludes giving a summary of the chapter and provides an introduction to Chapter 4.

3.2 Neutral Theory & Neutrality

Neutral Theory as discussed by Kimura [111], argues that mutation, not selection, is the main force in evolution. He describes how a mutation from one gene to another can be viewed as being *neutral* if it does not affect the phenotype, as the number of different genotypes which store genetic information is far greater than the number of phenotypes. This implies that the representation from genotype to phenotype must incorporate an element of redundancy so that neutral mutations are possible.

A mutation can be viewed as neutral if it changes the genotype but doesn't impact on the phenotype.

Previous work on evolution revolved around observed changes at a phenotypic level. Having access to molecular data, Kimura was able to experiment at a level not previously possible. As a result of examining the evolutionary rates of nucleotide substitutions, Kimura identified two types of gene mutations: *replacements of one nucleotide base for another* [112] and *structural changes consisting of deletions and insertions of one or more nucleotide's bases as well as transpositions and inversions of larger DNA segments* [112]. The results of Kimura's investigations led to the development of his theory of molecular evolution: *... a majority of nucleotide substitutions in the course of evolution must be as a result of random fixation of selectively neutral or nearly neutral mutants rather than positive Darwinian selection, and many of the enzyme polymorphisms are selectively neutral and maintained by the balance between mutational input and random extension* [112].

Although originally viewed as being anti-Darwinian, Kimura [112] stated that although natural selection is important in evolution, the number of DNA changes which are adapted in evolution are small, with the vast majority of mutations being phenotypically silent. Nei [139] argues that all evolution must be non-Darwinian and as phenotypic characters are under the control of DNA sequences, molecular and phenotypic evolution must take place under a similar non-Darwinian approach. Nei also argues that the majority of *morphological evolution* is as a result of neutral or nearly neutral mutation and therefore mutation can be considered the main force behind evolution at both the molecular level and the phenotypic level [139]. Following

Kimura, work by King and Dukes [113] describes how much of the evolution of proteins is down to neutral mutations and genetic drift.

Many studies focused on neutral theory, including Schuster [60, 163, 164], Huynen [100, 101], Reidys [149] and Shipman [167] all illustrating that by introducing redundant representations and thus neutral mutation, the connectivity between fitness landscapes can be altered. In other words, when a number of genotypes represent the same phenotype, they can be viewed as a *neutral set* and in turn alter the way in which a population explores the search space.

Neutral networks are normally viewed a collection of elements which are separated from one another by one Hamming distance. In this dissertation neutral networks are viewed as a set of points of similar fitness in the search space and starting from any given point within a set, an individual can reach any other individual in the set through one or more mutations, without having to leave the set [147]. Reidys et al. [150] suggest that neutral networks can be constructed by combining identical phenotype structures, so long as the structures exceeded a given threshold value. Their work was based on RNA secondary structures as they incorporate a high level of redundancy because of the existence of a greater number of sequences compared to structures. The mapping from RNA sequence to secondary structure can be described as a mapping from sequence space into shape space [164].

Schuster et al. [164, 163], examined mappings from an RNA sequence to an RNA secondary structure and concluded that in order to understand evolution, the genotype-phenotype mapping needs to be examined and that neutrality was a feature which can assist evolution. In Shipman [167] the author used random neutral walks.

He began with the random generation of a solution, then created a neighbourhood and selected a neutral neighbour that produces a distance increase from the starting point. The process is repeated until no increase is available and the walk stops. Fontana and Schuster [60] mapped an RNA sequence to a secondary RNA structure, but discovered that during the evolutionary process there existed periods where no adaptive progress took place, this was referred to as neutrality. Following this period of neutrality, there were periods where there was a significant increase in adaptive progress. This observation was important in terms of developing an understanding with regard to neutrality in the evolutionary process, as they also noted that there was activity, in that RNA shapes were changing during the level periods where there was no apparent change in adaptation.

Again looking at the mappings from RNA sequence to RNA secondary structures, Huynen [100] showed that although there are large amounts of neutrality present in the mapping, neutral paths exist, which allow for smooth exploration, with results indicating that the number of new structures discovered by neutral mutations increased linearly. Shipman et al. examined the properties of four redundant genotype-phenotype mappings: voting mapping, cursor based mapping, cellular automata mapping and random Boolean network mapping. In all cases, redundancy was found to be beneficial, meaning *movement on the resulting neutral networks allowed for the discovery of a larger number of phenotypes than would be the case for a direct encoding. Thus, the probability of entrapment at local optima when using these mappings would be reduced* [168].

Nimwegen et al. [143] illustrated the evolution of the population through neutral

networks, suggesting that migrating individuals do not randomly drift through them, rather the majority tend to remain highly connected, that is with neutral neighbours sharing the same level of fitness. This results in phenotypes which are quite resilient against mutation. Wagner [191] took a similar view and suggested that neutrality assisted in sheltering a system from the impact of mutation and argued that new adaptations, provided by neutrality, can assist the search. Wilke et al. [201] argued that neutrality provides robustness, with results indicating that genotypes closely connected by mutation, have high selection rates and that these genotypes tend to be located on flatter areas of the landscape.

3.3 Synthetic Neutrality

3.3.1 Types & Properties

In relation to neutrality in artificial systems, Weicker and Weicker [195] outlined four types of redundancy; coding based redundancy, representation based redundancy, conceptual redundancy and finally technical based redundancy. Coding based redundancy occurs when the size of the search space differs in relation to the size genotype space. The structure of the problem or the structure of the optimisation technique leads to representation based redundancy. Conceptual redundancy is created as a result of gene interaction and technical redundancy exists when a form of decoder is used to alter or repair the solution. Their research focused on conceptual and technical redundancy, and used a binary encoding with a decoder method and a diploid encoding as outlined by Dasgupta [42]. The study examines how neutrality changed

local optima to plateau points for both the decoder method and the diploid encoding. However, they also found the decoder method outperformed the diploid method.

Rothlauf and Goldberg [156] argued that the effects on evolutionary search of redundancy are dependent on the nature of the redundancy and identified a number of properties associated with redundant representations: *uniform*; *synonymous*; *locality* and *connectivity*, illustrating that the characteristics of the redundancy, which is dependent upon one of the outlined properties, has the greatest impact on search. The characteristics were defined as follows: uniform; a representation is considered uniform if all of the phenotypes are represented by the same number of genotypes. Synonymous; if the genotypes which represent the same phenotype are similar to one another, the representation is considered synonymous (and non-synonymous if they differ). Locality; if neighbouring genotypes correspond to neighbouring phenotypes, a representation is viewed as having a high locality of reference. In other words the locality of representation and synonymity are related; if neighbouring genotypes correspond to neighbouring phenotypes, then the representation has a high locality of representation and changes in the genotypes result in changes in the phenotype. Connectivity; if the number of phenotypes which are accessible by single-bit mutation from a given phenotype is high, then the connectivity of the representation is high [154]. Therefore, where a representation is viewed as synonymously redundant, the genotypes representing the same phenotype have the same properties and are next to one another in the mutation space. However, with non-synonymously redundant representations, as genotypes representing the same phenotype may be quite different from one another, leaving evolutionary search to operate like random search meaning

genetic operators will not function properly, thereby resulting in a lower performance [154].

With a redundant representation, a phenotype's phenotypic neighbourhood corresponds to the phenotypes which are reachable from the given phenotype by a single mutation of the genotype representing it [35]. Connectivity relates to the number of phenotypes which make up the phenotypic neighbourhood. With a non-redundant binary representation with a phenotype of length l , if each reachable phenotype from any another phenotype are all different, then the connectivity is l [35]. When defining neutrality, linear representations specify a set of phenotypes reachable from a given phenotype by a single mutation and all genotypes which represent the same phenotype reach the same set of phenotypes. However, when taking neutral theory into account, then the accumulation of neutral mutations should lead to new paths being located [35, 37]. A redundant representation which allows for the increase in connectivity should be synonymous as if a non-synonymously redundant representation is used, this will result in a random search which will have a negative impact on the efficiency of the search [154].

3.3.2 Neutrality in Artificial Systems

Shipman [167] found neutrality to be advantageous where neutral networks (defined by Harvey and Thompson [84] as points in a search space of equal fitness) are distributed over the search space with a high degree of connectivity between them. Shakelton [165] and Shipman [169] showed that neutrality could be introduced through the use of genotype-phenotype mappings. They used five different mappings

to illustrate this, *static random mapping*, *trivial voting mapping*, *standard voting mappings*, *cellular automata mapping* and a *random Boolean network (RBN)*. The static random mapping involved a genotype of length 30 being mapped to a phenotype of length 16. The mapping was created at initialisation and remained static. The trivial voting mapping, mapped three-bits from the genotype to one bit at the phenotype. The mapping involved calculating what the majority of the three bits at the genotypic level represented. For example, if two of the three bits were of value 1, then the genotype mapped onto value 1. Standard voting mapping varies from the previous mapping in that a set of bits can overlap at the genotypic level. This implies that when a mutation occurs many bits can be changed at the phenotypic level. Cellular automata mapping, which includes a truth table, takes in three Bits and determines the phenotypic representation. Finally, the RBN allows the three bits at the genotypic level to be at any loci on the chromosome, meaning that the position needs recording as well. The results obtained indicated that the amount of redundancy present was significant in relation to evolution. Knowles and Watson [114] argue that previous results on the benefits of including redundant encodings only show the increased connectivity of the redundant space and *is not a sound demonstration of the utility of the technique* [114]. They compared RBN encoding with direct encoding *given an appropriate choice of mutation rate* over NK landscapes [109], H-IFF [194] and Max-Sat [136], with results indicating that there was no longterm benefit in terms of the level of fitness achieved, when neutrality is included [114].

Yu and Miller [205] used a Boolean function landscape and showed neutrality to have a positive influence on evolutionary search. They examined explicit neutrality

and argued that when mutation occurs on active genes, it is adaptive, suggesting this comes from taking advantage of previous accumulated beneficial mutations within the population. They also observed that mutation on inactive genes is neutral and doesn't affect the fitness value associated with the genotype. It does however maintain diversity in the population and this can assist exploration. Yu and Miller [206] illustrated, using a Boolean function induction problem, that there exists a relationship between neutral mutations and beneficial search. Yu and Miller [207] also examined neutrality using a OneMax problem. The authors used a representation which included active and inactive genes. Each *node* in the genotype comprised of two genes. The first gene was viewed as an input link and the second as a function value. The function value is taken from a set $\{0, 1\}$, which either added 0 or 1 respectively. This representation allowed the measurement of neutral mutations using Hamming distance. Results showed that for the OneMax unimodal landscape, neutrality had a positive effect on search because it provides buffers to dampen the impact of destructive mutations.

Fonseca and Correia [57] used mathematical models to introduce two redundant representations. They found that the connection between phenotypes tended to increase as the number of redundant bits increased and that this trend in connectivity between phenotypes, exists even when low amounts of redundancy are present. This differed from the idea presented by Shipman [167], where in order to assist evolution, large amounts of neutrality must be present in the population. It has also been argued that landscapes with higher degrees of neutrality tend to have larger neutral networks [107]. Fonseca and Correia [58] developed a polygenic redundant representation which produced results similar to that of a non-redundant genotype-phenotype map. They

expanded the mapping to include pleiotropy, which allowed the connectivity to increase proportionally to the chromosome length with a non-neutral representation. Once they included error-control code theory to introduce neutrality they defined a group of redundant binary representations implementing various levels of neutrality, connectivity and locality.

Correia [35] extended this work using NK fitness landscapes and analysed the impact of redundancy and neutrality on the performance of an evolutionary strategy to gain a better understanding of how redundancy and neutrality affect evolutionary search. The performance of a (1+1)-ES was applied to NK Fitness landscapes and Markov chains were used to model its behaviour. A family of neutral network representations inspired by error-control codes, using a $NN(l, k)$ notation to signify a neutral representation and $NonNN(l, k)$ to denote a non-neutral representation. The $NonNN_z(l', k)$ family with the same phenotype neighbourhood was used for comparison in order to separate the neutrality from the effects of the neighbourhood common to both representations. The findings indicated that neutrality seems to affect the search less significantly than the phenotypic neighbourhood, resulting in delaying convergence in comparison with the non-neutral representation but not having a major effect on the algorithms long-term behaviour [35].

Further research by Fonseca and Correia [58] identify that linear encodings allow the specification of a set of phenotypes reachable from a particular genotype through single-bit mutations and that all genotypic representations of each phenotype are equivalent with regards to searching as they all reach the same group of phenotypes. But neutral theory [111] would suggest that changes to the connectivity between geno-

types and phenotypes are possible. Rothlauf [154] argued that when a synonymously redundant representation is used with a selectorecombinative GA, the connectivity between the phenotypes is not increased. However, Correia [35] argues that there are synonymously redundant representations that allow connectivity to be increased between phenotypes, when compared to non-redundant representations. Reporting that with their representation, connectivity increased with the amount of redundancy and that high connectivity can be found with very little redundancy. This finding contrasts with the idea that large amounts of neutrality are required to aid evolution [167].

Continuing research by Correia [36] illustrated that the influence of neutrality was separated from the effects of the neighbourhood common to both the neutral and non-neutral representations presented. Concluding that neutrality seems to have less of an effect on the search than that of the phenotypic neighbourhood, delaying the convergence of the algorithm in comparison with the non-neutral landscape, but significantly not changing the algorithm's long-term behaviour. The author also found that with the NK landscape presented, the representations did not present the values for connectivity, synonymity or locality expected from current literature.

This in turn may help in relation to finding a *demonstrably successful redundant representations in EC* [36], this research showed how the phenotypic neighborhood induced by a representation influences the performance of an evolutionary algorithm. Using a (1+1)-ES and the expected value of fitness, it was possible to conclude that the phenotypic neighborhood induced by the representation seems to dominate the behavior of the algorithm, affecting the search more significantly than neutrality.

In general, neutrality delayed the convergence of the strategy, but did not seem to significantly affect the long-term behavior of the algorithm. Among neutral representations in which phenotypic neighborhood contains the phenotypic neighborhood of the non-redundant representation, some of them presented a performance systematically superior than others. This observation suggests that there are representations that are more appropriate than others for the optimization of NK fitness landscapes.

3.3.3 Neutrality in GAs

Neutral representations have appeared in a number of GAs over the past number of years. As a general rule, the introduction of neutrality into GAs can be divided into two categories, neutrality through fitness landscapes and neutrality through genotype-phenotype mappings.

Neutrality through Fitness Landscape

The first category, fitness landscapes which introduce neutrality through landscapes such as Kauffman's NK landscape [109], Barnett's NK_p landscape [13], Newman and Engelhardt's NK_q landscape [140, 141]; and Beaudoin et al.'s ND landscape [18]. In Beaudoin's landscape N represents the length of the genome and D represents the *neutral degree distribution*, with the search space being divided into D neutral networks. The aim of this is to control and examine the distribution of neutrality, as they view this as being fundamental for evolution. By developing an ND deceptive landscape, the authors found that as the level of difficulty increases, the addition of neutrality was beneficial and when the level of difficulty decreased, neutrality didn't

assist in searching the search space.

Gerard et. al. [64] carried out comparisons of the NK , NK_P and NK_q fitness landscapes and observed that when neutrality is added there were a number of major differences associated with the distribution of neutral mutations. Smith et al. [171] examined the behaviour of the population during phases where a population's fitness remained relatively constant, the neutral phase. They argued that the presence of neutrality was not beneficial for evolutionary search. Smith et al. [172, 173] concluded that the neutral phases were as a result of the presence of neutrality. Smith et al. [170, 171] also proposed the use of fitness evolutionary portraits, obtained by calculating the average evolvability of a population of equal fitness. Katada et al. [108, 106] describe the relationship between evolution and neutrality in terms of two periods, the transient period and the equilibrium period. During the transient period fitness levels improve whereas during the equilibrium period, fitness levels remain constant. They noted that the equilibrium period lasted longer than the transient period. Van Nimwegan et al. [142] also examined the equilibrium period (stasis) and the transient period where rapid change is observed. These are but a few of the examples of the examination of neutrality through the use of landscapes. Others examples include Lobo [121] Smith [172] and Yu [205, 207, 206].

Neutrality through Genotype-phenotype Mappings

Moving to the second category, the introduction of neutrality through genotype-phenotype mappings, which is the approach taken in this dissertation towards the introduction of neutrality, Lewontin observed that *If we take Darwin's view that evo-*

lution is the conversion of variation between individuals into variation between populations and species in time and space, then an essential ingredient in the study of evolution is the study of the origin and dynamics of genetic variation within populations [119]. Lewontin [119] identified two distance spaces, a genotype space G and a phenotype space P and illustrated the representation of the structure of population genetics as:

$$G_1 \xrightarrow{T_1} P_1 \xrightarrow{T_2} P_2 \xrightarrow{T_3} G_2 \xrightarrow{T_4} G'_1 \xrightarrow{T_1} \dots$$

Which provides a set of rules that maps a population of genotypes (G_1) to a phenotype space (P_1), for selection to occur [56]. Another set of rules map the resulting population (P_2) back to genotype space (G_2) where Mendelian genetics can predict the next generation of genotypes [56]. The four laws associated with transformation are; T_1 epigenetic; T_2 natural selection; T_3 genotypic survival and T_4 mutation [56]. Lewontin points out that population genetic theory maps a set of genotypes into a set of phenotypes, provides a transformation in the phenotype space, then map these new phenotypes back into genotypes, where a final transformation occurs to produce the genotypic array in the next generation [119].

Another aspect of the mapping between the genotype and the phenotype is the presence of *pleiotropy* and *polygenic* inheritance. With *pleiotropy* a single gene can impact on a number of phenotypic traits, *polygene*, on the other hand, a single phenotypic trait may be determined by the interaction of many genes [119]. If we consider that the operators used in association with GAs are applied at the genotypic level [67], then when viewing a GA in terms of Lewontin's mappings [119], GAs operate from genotype to genotype and new genotypes are created by implementing an inter-

pretation of biologically inspired genetic mechanisms [119]. The fitness of a genotype is based on the evaluation of its phenotype.

Evolution explores the phenotypic search space through mutation, recombination and selection, with *exploration being governed by the variational topology of possible mutations on phenotype space* [184]. When discussing variational topology, that is the effects of changes in spacial relations amongst individuals, we need to examine both genotypic variational topology and phenotypic variational topology. Nature uses a complex genotype-phenotype map to advance a relatively simple genotype space variational topology to an extremely complex phenotypic variational topology. Toussaint [183] argues that the phenotype space is what should in fact be viewed as the search space for evolution rather than the genotype space. With this in mind, the genotype-phenotype map is the key to understanding phenotypic variability and it also allows us to gain an insight into how evolution can adapt the search on the phenotype space.

The genotype is a collection of items which are inherited and define the possible distribution of offspring. The phenotype on the other hand, comprises all of the phenotypic traits of an individual that contribute to the overall fitness of that individual. The genotype-phenotype map sets out a blueprint for moving from genotype to phenotype. By viewing a phenotype as the combination of all of the traits of an organism and the genotype as the genetic encoding of this phenotype, then we will also assume that there exists a genotypic neighbourhood and a phenotypic neighbourhood. Two genotypes can be considered equivalent if they represent the same phenotype, however their variational topologies can differ [176]. The genotype-phenotype map “*induces a*

variational topology on the phenotype space depending on the topology of the genotype space” [183].

If we allow for a surjective genotype-phenotype map, then the same phenotype can be encoded by a number of genotypes. One major advantage of allowing a many-to-one non-trivial genotype-phenotype map is that it enables the changing of the genotype representation without changing the phenotype. This can be achieved through *neutral mutations* which allow for changes in the phenotypic neighbourhood and impact on the search. In other words, neutrality enables phenotypic variability to adapt and vary, whereby the genotype-phenotype map itself doesn’t need to evolve, neutrality in the phenotypic variational space allows changes to occur in the gene interaction which introduces a different phenotypic variability. Therefore, by introducing neutrality in a non-trivial mapping, the phenotypic neighbourhood depends on the genotypic representation of the phenotype and major alterations of the genotypic representation mean changes in the phenotypic neighbourhood [183]. “*The notion of fitness landscapes is sometimes also used to describe a fitness function over the phenotype space in the case of a non-trivial genotype-phenotype mapping*” [183]. But by introducing a fixed (non-trivial) genotype-phenotype map you can still achieve adaptable exploration through “*neutral variations in the genotype space*” [183].

Trivial neutrality can be defined as a situation where “*the evolution of phenotypes can be understood (i.e., modelled) without referring at all to genotypes, in particular, neutral traits are completely irrelevant for the evolution of phenotypes*” [183]. In other words with a fixed genotype-phenotype mapping, if their phenotypic outcomes, that is their *phenotypic mutation distributions, are constant in each neutral set* [147].

Therefore with trivial neutrality, neutral traits do not have an impact on phenotypic evolution [147]. In a non-trivial mapping a single gene is associated with a number of features in the phenotype, when this gene mutates the features vary in accordance [183]. This means that with neutrality which is non-trivial, different genotypes within *a neutral set induce different phenotypic distributions* [147]. In Toussaint's [183] view this implies that selection between genotypes of a similar fitness is *similar to the selection of strategy parameters in self-adaptive EAs* [147].

For a genotype-phenotype map to be non-trivial there are a number of conditions which must be met. Firstly, there is a surjective relationship between genotype and phenotype and secondly, for a number of the phenotypes, the phenotypic variability generated depends on the genotype from which it has emerged and the genotype contains neutral traits, which implies that you can have different genetic representations for the phenotype [182].

Toussaint's work [184] using *trivial neutrality*, phenotypic mutation distributions, are constant over individual neutral sets, indicated that neutral traits have no impact on phenotypic neutrality [147]. Results showed that trivial neutrality is a necessary condition for compatibility with a phenotypic projection of a mutation-selection GA i.e. *whether one or another representative of a neutral set is present in a population does not influence the evolution* [147] of a phenotype. What this means is that with trivial neutrality, neutral traits have no impact on phenotypic evolution. On the other hand, with non-trivial neutrality, different genotypes in a neutral set allow for different phenotypic distributions [147]. Toussaint viewed this as having similarities with the selection of strategy parameters in self-adaptive evolutionary algorithms, as

there exists a selection between equivalent genotypes when a non-trivial neutrality is present. In other words he viewed it as fundamental to the mechanics of evolving genetic representations [184]. Toussaint and Igel [185] argued that approaches to self-adaptation in evolutionary algorithms can be viewed as an example of the benefits of neutrality. As each chromosome has a mutation rate associated with it, chromosomes may be part of the same neutral network, as a result of the presence of neutrality, and have different mutation rates and evolution can choose between these in a self-adaptive way. The significance of this is that there exists a variety of mutation rates or distributions within the population, which can evolve. By extension, theoretical work on self-adaptation i.e. [53, 4, 187, 188] etc. may be viewed as examining the evolution of neutral traits or as explicit examples of the benefit of neutrality [147].

Many authors suggest that the introduction of redundant representation into a genetic algorithm introduces an additional source of diversity within the population. This in turn implies that *implicit parallelism* occurs for an extended period of time. Cohoon et al. [33] draw on the idea of *punctuated equilibria*, which incorporates the principles of *allopatric speciation* (rapid evolution of a new species after a small subset of the population become segregated in a new environment) and *stasis* (where an isolated population, over time, stabilises). Different environments are created, by allowing fitness to be accessed relative to a local population. The *species* that develop in each environment can be viewed as equivalence classes. In this way members of different environments can be exchanged and thereby introduce a new level of competition into an environment. By having equivalent classes a many-to-one representation exists, which proved beneficial to the search using a modified GA. Gould

[77] explains punctuated equilibria by claiming that the idea that all change is gradual and continuous results from interpreting changes as mainly adaptations, with most retained changes being neutral. *The accumulation of neutral changes may produce pre-adaptation phenomena because retained neutral characters may later turn out to be adapted to new circumstances. Pre-adaptations may in turn explain punctuated equilibria* [27]. Therefore when a *neutral character becomes useful through a new additional change, the fitness and/or behaviour may change suddenly* [27]. Cangelosi et al. [27] introduced the biological concepts of cell division and migration in addition to axonal growth and branching into the representation, building on work by Miglino et al. [130] and found that changes in early stages of cell division and migration had dramatic effect on the phenotype, with changes in later stages having less of an effect. The results also illustrated that *the effect of a single mutation on the phenotype (the network) can be much greater if the mapping from genotype to phenotype is complex and many-to-many rather than simply one-to-one* [27].

Banzaf proposed an application of a genotype-phenotype map which separated the genotype search space from the phenotype solution space [10]. As the phenotype is an expression of the information contained in the genotype, Banzaf's mapping included an interpretation of transcription and translation and guaranteed the feasibility of the phenotype, which resulted in multiple genotypes representing the same phenotype [10]. This allowed for neutrality in the representation, which assisted in maintaining diversity. Also as described by Elgin [54], random drift caused by neutral variants increases population diversity by expanding the population distribution so as to help find an escape route from local traps, particularly in high dimensional

spaces. Translation was also adopted by Ashlock through the use of a translation table into a self-avoiding walk test problem, but yielded no significant improvement in performance [5].

Lehre and Haddow [117, 118] developed a mapping which converted elements into two symbols, based on RNA secondary folding. The authors identified two types of neutrality: *step k-neutrality*, which relates to neutrality created during the mapping and *remaining neutrality*, which refers to neutrality which is visible after mapping has taken place. Their results indicate that with *step k-neutrality*, the quantity of neutrality present varies at the genotypic level and for phenotypes which were complex, the quantity of neutrality is relatively low. Results also showed that when phenotypes were less complicated, the quantity of neutrality present was high. They also illustrated that the size of neutral networks varied from small to large. The approach of using mappings was extended by Ebner et al. [51, 52], who took two of the mappings, cellular automata and RBN with what they referred to as *phenotype-species mapping*. They outlined how high levels of mutation could be sustained by having neutral networks present. They also identified that neutral networks assist in maintaining diversity in the population, which may be advantageous in a changing environment.

Chow et al. [29, 30] outlined a diploid chromosome, a data chromosome and a mapping chromosome. Both chromosomes undergo an evolutionary process, the data chromosome is altered by crossover and mutation and for the mapping chromosome bits are altered by using a permutation operator. This type of mapping introduces neutrality as all the bits of the data chromosome may not take part in

genotype-phenotype mapping. Using trap functions [47], Chow illustrated that there were advantages associated with neutrality. Ronald et al. [152] examined the use of redundant representation and describe how representative encodings can introduce *cloaked duplicates*, where a number of different genotypes (also referred to as *isomorphic genotypes*) encode for the same phenotype. The authors point out that with GAs and scheduling problems, a degree of redundancy exists in the encodings and use the Travelling Salesperson Problem to illustrate the presence of *cyclic-shift isomorphism* and *inverted-ordering isomorphism*. Because of the presence of these isomorphisms, an Edge Recombination Operator (ERO) (as found in [200]) and an inversion operator were used for crossover and mutation. A process of isomorphic normalisation was used to reduce redundancy and in turn, the size of the search space. The paper argued that by eliminating redundancy, through normalisation and duplicate removal was beneficial.

Galván-López and Poli [63] introduced a neutral network of constant fitness, equally distributed across the entire search space. They used a mutation only binary GA and added neutrality by including an extra bit to the representation. A 1 indicated that the genotype is on the neutral network and a 0 indicated that the genotypes fitness is determined by the coding bit as would be usual, dividing the search space into two equally sized spaces, the neutral layer and the normal layer. The authors analysed population flows; compared the percentages of successful searches and the difficulty of the landscape using fitness distance correlation. A OneMax problem and a deceptive trap function were used as test problems. The authors argued that neutrality may be beneficial in some cases. However this comes at a cost, as there is

extra computational overhead required due to the increase in the search space. Which means that any additional benefit may be outweighed by the additional overhead.

The unimodal landscape used, showed little sign of any benefit of including neutrality. There may however, be a benefit of including neutrality when the population is initialised in the wrong area of the search space. The authors also tracked the origin of a sample point and gathered data on the population flows, as each individual had only one parent, from one area to another. The population flow data indicated that the majority of off-spring in an area come from parents already in the area. The authors illustrated that there is a difference in the behaviour of a genetic algorithm depending on whether a neutral network has a low or a high fitness value. They argue that neutrality should be beneficial only when it *modifies the search basis of an algorithm-problem pair* [62] which improves the likelihood of sampling the global optimum. If the reverse happens and the introduction of neutrality makes it more difficult to reach to global optimum, then neutrality will be of little benefit [62]. Doerr et al. [49] carried out a run-time analysis of the effects of incorporating neutrality. They used the same model of neutrality and problems as found in [63, 62, 61] and show that there was no significant advantage of using neutrality over a OneMax landscape or a deceptive problem landscape with a single global optimum and a single local optima. They did view neutrality as being possibly beneficial over deceptive functions with a single global optima and two local optima, as the run-time was reduced from exponential to polynomial or when the *number of search points being nearly optimal is at least a polynomial fraction of the search space* [49]. The results illustrate how neutrality may be beneficial, depending on the amount of neutrality present and the

type of landscape.

3.4 Chapter Summary

To summarise, this chapter reviewed the literature associated with Neutral theory; neutrality from a biological context; and synthetic neutrality as implemented in artificial systems. As neutrality can be defined as a situation where following a mutation one genotype changes to another genotype, but both genotypes represent the same phenotype [111], this implies that as neutrality is introduced, the solution space increases without increasing the genotype space. The main focus of most of the research to-date relates to the evolvability of a population and examines the ability of variations to assist the search by including redundant representations, with research examining the impact of being able to change the genotype without changing the phenotype. As this dissertation examines the impact of implementing Neutral theory in a GA, the chapter focused primarily on the introduction of synthetic neutrality into artificial systems and specifically on neutrality in GAs, where neutrality is introduced through either fitness landscapes or through the genotype-phenotype mapping.

Having examined past research one of the most interesting questions lies in the development of a non-trivial representation and examining the impact of they type of neutrality produced by the representation on variation and variability. In developing a tunable neutral representation, which maintains variation and alters variability, it is planned to use this framework to empirically analyse the impact of neutrality over various landscapes using a GA which incorporates the modular based, fixed, non-trivial genotype-phenotype representation. The next chapter outlines the design of

the multi-layered GA (MGA), and describes how neutrality is introduced into a GA's genotype-phenotype mapping, through an implementation of the biological processes of *transcription* and *translation*, discussing dynamic activation and its impact on connectivity. The chapter also outlines the missense mutation operator and how it functions within the layered representation.

Chapter 4

Multi-Layered GA (MGA) Design

4.1 Introduction

In designing the proposed representation, the aim is to develop a representation for the MGA which is modular in construction, maintains variation, increases variability and governs the evolutionary trajectory. The other design constraints are: that the layered mapping of the representation contains a binary haploid primary structure which allows the use of standard well understood, variational operators in a problem independent manner; that multiple genotypes will represent the same phenotype [10]. This chapter outlines the motivation and design of the proposed MGA representation, which includes a many-to-one genotype-phenotype mapping. The proposed mapping is modular in design and incorporates a fixed, non-trivial mapping. The chapter is laid out as follows: Section 4.2 outlines the motivation behind the MGA design. Section 4.3 gives an overview of the MGA, describing its design and Section 4.4 outlines the implemented interpretation of the transcription and translation phases,

allowing the two step expression of the genes. The chapter also describes the missense mutation operator and outlines the arity of representations used in the dissertation, giving a worked example of a 6-3-1 representation. Section 4.5 presents an overview of the workings of the genotype-phenotype mapping, giving a detailed and worked description of the mapping from the genotype through the processes of transcriptions and translation and on to create the *phenes* and eventually the phenotype. A phene can be used to describe a single unit of a phenotype, where a phenotype is the manifestation of it's phenes. Finally, Section 4.6 concludes the chapter.

4.2 Motivating the MGA

Rather than introducing an initial *Big Bang* of diversity into an unusually large population to decrease the probability of premature convergence [14], a GA which through its representation contains neutrality, maintains diversity and prevents premature convergence may offer a useful search strategy over difficult landscapes. In relation to *organismic design*, modularity is a common feature [192] and can be interpreted as a genotype-phenotype map where pleiotropy (a change in a gene can influence a number of traits) mainly exists *amongst characters that are part of a single function complex* [192]. By adopting this approach, evolvability can improve as the levels of interference between the *adaption of different functions* is reduced [192]. The genotype-phenotype map included in the MGA incorporates modularity and the idea of characters being part of a single function complex, whilst minimising pleiotropy.

The variational properties of the phenotype are a level of phenomena distinct from

phenotypic adaptation; they are subject to distinct evolutionary dynamics [192]. When exploring the phenotypic space, it is critical to gain an understanding of the variational topology in trying to determine the shape of the landscape [184]. Many evolutionary algorithms are created using a fixed variation topology. However, in nature phenotypic variation landscapes are not fixed. These non-fixed phenotypic variation landscapes can be referred to as a non-trivial in terms of their genotype-phenotype map [184]. A non-trivial genotype-phenotype map can be viewed as having the following characteristics: firstly, a phenotype can be encoded by many genotypes and secondly, the phenotypic variability of a number of phenotypes will depend on their genotype [182]. Therefore, every genotype can be described as a combination of a phenotype and any kind of neutral traits which exist within the genotype, where different neutral traits associated with the same phenotype give a different representation of the phenotype in question [182].

Many authors have developed GAs which attempt to introduce interpretations of biological processes including [204, 145, 26, 28, 116]. For GAs, the primary structure (genetic representation) is the genotype and the secondary structure is the phenotype. Also in GAs, the variational topology is traditionally fixed, whereas in nature phenotypic variational topology is not fixed. However, as pointed out by Toussaint in [184], the introduction of an adaptable exploration is possible by creating a fixed but non-trivial genotype-phenotype map, which is achieved by having neutral variations in the genotype space. The choice of genetic representation, the primary structure, used to represent a neutral set which encodes a phenotype will impact on the phenotypic variability. The outcome of this is that the evolutionary path is governed by

phenotypic variational topology.

4.2.1 Phenotypic Variability & Neighbourhood Equivalence

The idea of including redundancy or neutrality into the genetic representation introduces the notion of phenotypic equivalence, as neighbourhoods of genotypes can be formed based on the idea of equivalent classes. In other words, as individual alleles only appear for natural selection when they are expressed phenotypically, alleles which are phenotypically similar can be clustered into phenotypic equivalence classes. However, this idea of comparison through the use of equivalent classes can be found throughout the literature, and measuring evolutionary activity at the level of genotypes or equivalent classes of genotypes is also straightforward [148]. But carrying out analysis at the level of alleles can be more difficult, particularly when the genotype-phenotype map is context sensitive, that is when you have epistasis or polygene, where many genes together can affect a particular trait or character [148]. Therefore, the interpretation of an alleles' evolutionary activity is easier if it has a phenotypic function which is clearly identifiable. The analysis of evolutionary activity becomes more difficult as the level of context sensitivity within the genotype-phenotype map increases.

To understand the evolutionary process it is critical that we understand the accessibility of phenotypes from genotypes and the impact of neutrality in relation to this mapping from genotype to phenotype. In order to gain an understanding of phenotypic variability we need to examine the evolutionary path from the genotypic space to the phenotypic space. These respective spaces (genotypic and phenotypic) can be

viewed as topological spaces which in turn are broken into sets of neighbourhoods. In the genotype topological space, particular genotypes are grouped into neighbourhoods based upon closeness to particular phenotypes. This closeness is not based on similarity, but on the accessibility of the phenotype to the genotype topological space and ignores any associated fitness values. The closeness can be determined by the number of mutations away a particular genotype is from a phenotype [59].

The phenotype space can be viewed as the phenotypic topological space which can also contains a set of neighbourhoods. However, the difficulty here is that phenotypes, by their nature are not altered by physical processes such as mutation. To overcome this, neighbourhoods in the phenotypic topological space can be classified by closeness in terms of a continuous evolutionary path. Therefore, by examining the statistical frequency by which a mutation on a particular genotype yields a phenotype, we can define a notion of closeness in the phenotypic topological space [59]. Neutrality and neutral genetic drift in the mapping process is important in defining the phenotypic topology, and the variation in phenotypic topology. Because of the presence of neutrality, phenotypes can be used to mean similarity between genotypes. Therefore, the closeness of phenotype neighbourhoods depends on the statistics of neighbourhood closeness among genotype neighbourhoods [60].

4.3 MGA Design

The Multi-layered Mapping GA (MGA), employs a modular approach to the creation of a genotype-phenotype map, which moves from genes to phenes by using a basic interpretation of the biological processes of transcription and translation. It is

these processes that create the neutral networks and define the variation and variability. At a very basic level, the biological process of transcription involves the copying of information stored in DNA into an RNA molecule, which is complementary to one strand of the DNA. The process of translation then converts the RNA, using a predefined translation table, to manufacture proteins by joining amino acids. These proteins can be viewed as a manifestation of the genetic code contained within DNA and act as organic catalysts in anatomy.

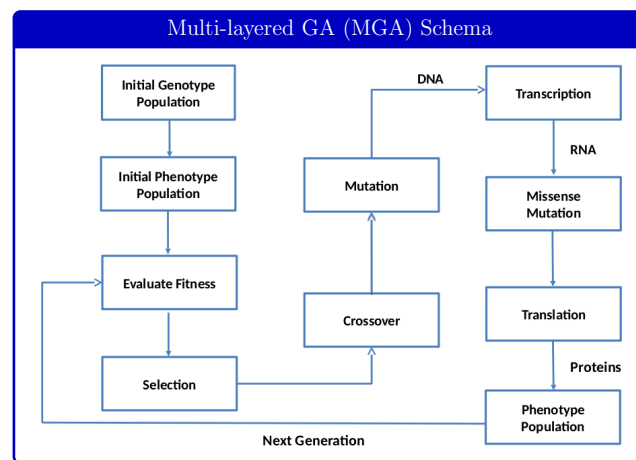
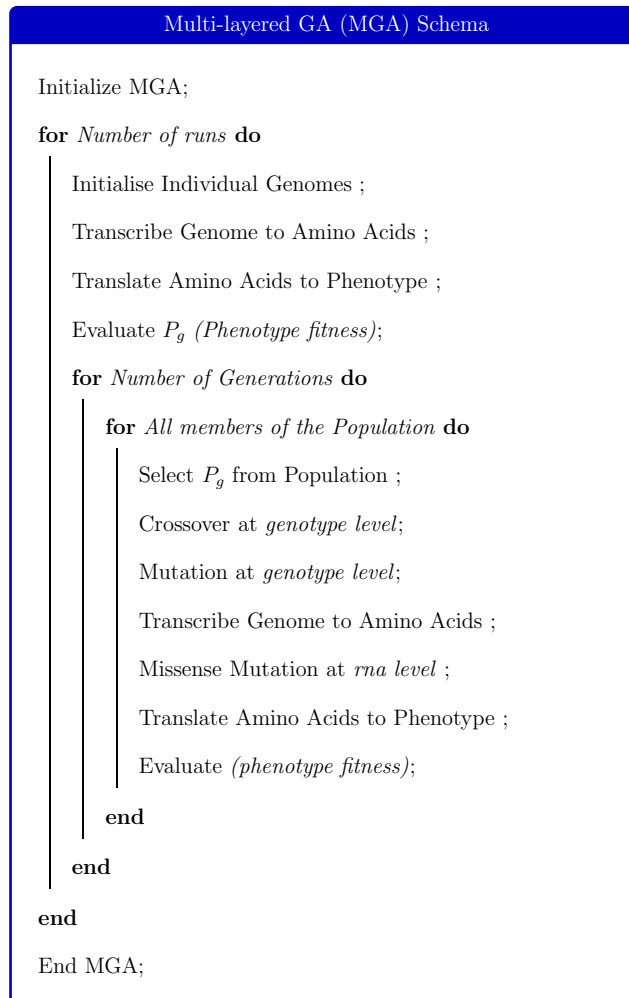


Figure 4.1: Overview of Multi-layered GA (MGA)

The schema for the MGA is depicted in Figure 4.1 and the mapping used in this dissertation produces two kinds of phenes, '0' or '1', each made up of combinations of a number of characters (A, C, G and U) representing amino acids. The number of characters used to represent a phene and the number and type of phene can vary and are set at initialisation. These phenes are then combined to create a phenotype which is evaluated for fitness. And by a phenotype I mean a collection of phenes,

with each phene incorporating a number of characters made up from genes, within the metaphor of EAs.



Algorithm 1: MGA Pseudocode

This approach enables changes in the phenotypic neighbourhood, where phenes, which are genetically determined characteristics or traits created through various combinations of *amino acids* based on translation tables. That is, if the translation tables change so too do the phenotypic neighbourhoods, as the translation phase

determines the phenotypic neighbourhoods. Another feature of the MGA is the ability to introduce variation operators at levels other than the genotypic level [10]. This offers the opportunity to include interpretations of operators [56], such as missense mutation, which exist in biology and operate at the RNA level. The pseudocode is outlined in Algorithm 1.

4.4 MGA Representation

The MGA employs a modular approach to the creation of a genotype-phenotype map, which moves from genes to phenes using a basic interpretation of transcription and translation. With the multi-layered mapping of the MGA, as you progress through the layers the level of redundancy alters. In the transcription layer, which maps $|\phi_g|$ (the genotype space) to $|\phi_d|$ (the DNA space) which is mapped to $|\phi_r|$ (the RNA space) there is no redundancy, however redundancy exists during translation. The MGA genotype-phenotype mapping comprises of two main phases, the transcription phase and the translation phase, both of which are outlined below.

4.4.1 Transcription Phase

The genotype-phenotype map enables the encoding of a phenotype of length l by a genotype of length lc , with c being the number of bits required to represent each of the k symbols, with k being the size of the alphabet being used by the MGA to represent a phene, in this dissertation ‘0’ and ‘1’. As all the genotypes bits have an effect on the evolution of the phenotype there is no redundancy and neutrality is not trivial i.e. the phenotype can be encoded by many genotypes and phenotypic

variability of a number of phenotypes will depend on their genotype [182].

The MGA also offers the ability to alter the alphabet size and to implement a neighbourhood map which allow for different types of neutral traits represented in the genotype to represent the same phenotype. Beginning with a binary string, which represents the genotype, the MGA maps pairs of binary bits into one of four characters A, C, G or T. These characters represent the first phase of the transcription where a *template* strand is created. Following this, the template strand maps onto a *coding* strand and the final stage of the transcription phase maps the coding strand onto RNA. Using a four letter alphabet (A, C, G and U), the RNA sequence is complementary to that of the DNA template strand and therefore is the same sequence as the DNA coding strand, with uracil (U) in place of thymine (T). Table 4.1 outlines the mapping process used by the MGA to implement the transcription stage. Each of the mapping layers offer the opportunity to introduce variation operators.

Transcription Process		
Template Strand Map	Coding Strand Map	RNA Map
00 → A	A → T	T → U
01 → C	C → G	G → G
10 → G	G → C	C → C
11 → T	T → A	A → A

Table 4.1: Transcription Stage - Template Map, Coding Map & RNA Map

4.4.2 Translation Phase

Once the transcription phase is completed and the alphabetic characters have been combined, the translation stage uses the information and maps it against the translation tables created by the MGA to examine neighbourhood equivalence and manufacture a phene, the MGA can also adjust the number of bits required to make up a phene and this can be set at initialisation. The size of the translation table is determined by the representation chosen. For a 4-bit MGA representation, a translation table of 2^4 is used; with a 6-bit MGA representation a 2^6 translation table is needed and with an 8-bit MGA representation a 2^8 translation table is created. The size of the translation represents the granularity of neutrality which exists in the representation.

6-3-1 Representation

In order to describe the MGA representation, the steps required to create a 6-3-1 representation (illustrated in Figure 4.4) are outlined below. In this example each phene is made up from 3 characters and each character is represented by 2 binary bits. Therefore, in the 6-3-1 MGA representation $c = 6$ as 6 bits are required to represent a single phene, '0' or '1'. As the value of c changes, so too does the amount of neutrality present in the representation and the accessibility of phenotypes from genotypes. Once the level of c and the number of phenes required has been decided, initialisation begins by creating the translation tables to represent each phene.

The creation of a set of translation tables requires integer values representing each phene, to be translated into binary strings is illustrated in Table 4.2, with Table 4.2a showing the 6-bit strings representing phene '0' and Table 4.2b the 6-bit

Phene '0'				Phene '1'			
000000	000010	000100	000110	000001	000011	000101	000111
001000	001010	001100	001110	001001	001011	001101	001111
010000	010010	010100	010110	010001	010011	010101	010111
011000	011010	011100	011110	011001	011011	011101	011111
100000	100010	100100	100110	100001	100011	100101	100111
101000	101010	101100	101110	101001	101011	101101	101111
110000	110010	110100	110110	110001	110011	110101	110111
111000	111010	111100	111110	111001	111011	111101	111111

(a) 6-3-1 Binary Representation Phene '0' (b) 6-3-1 Binary Representation Phene '1'

Table 4.2: Binary Representation for Phenes '0' & '1'

strings representing phene '1'. Once binary tables are created they are then mapped to translation tables containing combinations of characters. In this representation, 3 character groupings are used to represent each phene. The translation tables created for the 6-3-1 representation are shown in Table 4.3, with Table 4.3a containing the character combinations required for phene '0' and Table 4.3b, the characters required for phene '1'. These tables indicate that each of the phenes are represented by a combination of 3 characters.

Phene '0'				Phene '1'			
AAA	AAG	ACA	ACG	AAC	AAU	ACC	ACU
AGA	AGG	AUA	AUG	AGC	AGU	AUC	AUU
CAA	CAG	CCA	CCG	CAC	CAU	CCC	CCU
CGA	CGG	CUA	CUG	CGC	CGU	CUC	CUU
GAA	GAG	GCA	GCG	GAC	GAU	GCC	GCU
GGA	GGG	GUA	GUG	GGC	GGU	GUC	GUU
UAA	UAG	UCA	UCG	UAC	UAU	UCC	UCU
UGA	UGG	UUA	UUG	UGC	UGU	UUC	UUU

(a) Translation Table Phene '0'

(b) Translation Table Phene '1'

Table 4.3: Character Representation for Phenes '0' & '1'

Dynamic Activation

After information located in DNA is copied, through *transcription*, to RNA, *translation* takes place where RNA is used to assemble proteins. As much of the DNA does not encode proteins and instead encodes various types of functional RNAs [15], the proposed representation uses predefined DNA locations to code for a basic interpretation of different functional outputs. This allows the possibility of loci on a genotype, where given a certain context are *functionless* (*mutation makes no difference*) [84], but combined with a value elsewhere on the genotype may become important [84]

Dynamic Activation				
Genotype	Template Strand	Coding Strand	RNA Strand	Phenotype
0000 *** **0000	*AA *** *AA	*TT *** *TT	*UU *** *UU	0*0
1111 *** **1111	*TT *** *TT	*AA *** *AA	*AA *** *AA	1*1

Table 4.4: MGA 6-3-1 Dynamic Activation Pairings

The mappings shown in Table 4.4, identify the predefined locations and relate to the idea of dynamic activation [193, 15], which results in an adaptive change that produces the opposite outcome to that outlined in the translation table when pairs of characters are located on the predefined loci, match the pairings shown in Table 4.4. The ‘*’ represents don’t care in relation to characters present.

The dynamic activation encodings assist in increasing connectivity and alters the functional outcome of phenotypes. The process involves locating the presence of characters at predefined locations and if required, altering the nearest phenone value associated with the loci as per the mappings in Table 4.4 rather than per the trans-

lation tables. Dynamic activation enhances the non-trivial and pleiotropic nature of the representation and implies that a single gene is associated with a number of features in the phenotype and when the gene is mutated the features vary accordingly. A worked example of this is presented in Table 4.5.

Variability

As mentioned with a redundant representation, a phenotype's phenotypic neighbourhood corresponds to the phenotypes which are reachable from the given phenotype by a single mutation of the genotype representing it [35]. Connectivity relates to the number of phenotypes which make up the phenotypic neighbourhood of a phenotype. With a non-redundant binary representation with a phenotype of length l and each reachable phenotype from any another phenotype are all different, then the connectivity is l [35].

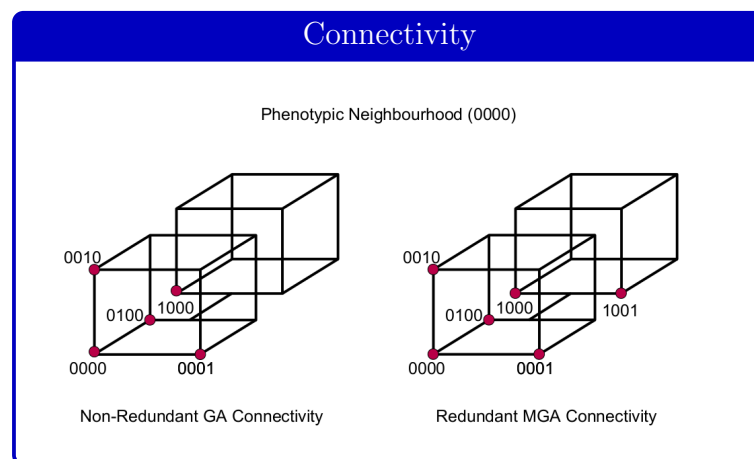


Figure 4.2: Non-Redundant & Redundant Phenotypic Connectivity

When defining neutrality, linear representations specify a set of phenotypes reachable from a given phenotype by a single mutation and all genotypes which represent the same phenotype reach the same set of phenotypes. However, when taking neutral theory into account, then the accumulation of neutral mutations should lead to new paths being located [35]. A mutation within the MGA model is viewed as being neutral when the resulting phenotype remains unchanged. Figure 4.2 compares the variability associated with a non-redundant GA representation and the proposed MGA redundant representation. The example illustrates the change in variability and shows a number of genotypes representing the 0000 phenotype can increase the connectivity for a phenotypic neighbourhood from 4, for the non-redundant representation to 5 through single bit adaptive mutation of the genotypes representing the 0000 phenotype. This change in connectivity is only applicable in situations where dynamic activation takes place. This is because dynamic activation, assists in minimising pleiotropy and in making the representation non-trivial, as the variability of a number of phenotypes will depend on their genotypes.

MGA Connectivity				
Polygenic & Pleiotropic Genotype	Template Strand	Coding Strand	RNA Strand	Phenotype
0111**11****11**1111	*CT **T **T *TT	*GA **A **A *AA	*GA **A **A *AA	0000
1111**11****11**1111	*TT **T **T *TT	*AA **A **A *AA	*AA **A **A *AA	1001
0110**11****11**1111	*CG **T **T *TT	*GC **A **A *AA	*GC **A **A *AA	1000
0111**10****11**1111	*CT **G **T *TT	*GA **C **A *AA	*GA **C **A *AA	0100
0111**11****10**1111	*CT **T **G *TT	*GA **A **C *AA	*GA **A **C *AA	0010
0111**11****11**1110	*CT **T **T *TC	*GA **A **A *AC	*GA **A **A *AC	0001

Table 4.5: Reachable Phenotypes through single Bit Adaptive Mutation

SGA Connectivity	
Genotype	Phenotype
0000	0000
1000	1000
0100	0100
0010	0010
0001	0001

Table 4.6: Reachable Phenotypes through single Bit Adaptive Mutation

The design of the MGA's primary structure is outlined in Table 4.5, which illustrates the impact of a single-bit mutations for a genotype representing the 0000 phenotype for the redundant 6-bit MGA representation and Table 4.6 for the non-redundant SGA representation. From the Table 4.5 each of the MGA's genotype contains '*' characters representing silent mutations (cold zones) and the '0' and '1' gene represent the possibility of adaptive mutations (hot zones). The hot zones represent a possibility of producing an adaptive mutation as not every mutation will lead to a different phenotypic trait [139]. The adaptive mutations are determined by the sets of characters and a translation table to which they belong. The predefined locations for dynamic activation are identified in blue and dynamic activation takes place in this example when two pairs of *A* characters are located on the RNA string, identified in red on the genotype.

4.4.3 Missense Mutation Operator

Missense mutation changes a single base pair which causes the substitution of a different amino acid in a resulting protein. The missense mutation operator is in-

terpreted and implemented as follows: once the processes of *transcription* has taken place, if a missense mutation occurs then one of the RNA bases is flipped to another, with the mapping used for missense mutation illustrated in Table 4.7, showing that each character can be mutated to any other character. Missense mutation is capable of producing silent and adaptive mutations depending on the position of the loci involved. Any changes introduced by missense mutation are reflected back on the genotype as well as altering the RNA string. Following missense mutation, the translation phase takes place. It should be noted that like traditional point mutation operators, the probability of a missense mutation taking place is on each single codon.

Missense Mutation Mapping			
A → C	Adaptive/Silent Mutation	C → A	Adaptive/Silent Mutation
A → G	Adaptive/Silent Mutation	C → G	Adaptive/Silent Mutation
A → U	Adaptive/Silent Mutation	C → U	Adaptive/Silent Mutation
G → A	Adaptive/Silent Mutation	U → A	Adaptive/Silent Mutation
G → C	Adaptive/Silent Mutation	U → C	Adaptive/Silent Mutation
G → U	Adaptive/Silent Mutation	U → G	Adaptive/Silent Mutation

Table 4.7: Missense Mutation Mapping

4.4.4 MGA Arity

The arity of the MGA, which impacts on the accessibility of phenotypes from genotypes, is decided by the user in advance of initialisation. In this dissertation three granularities of representation are used. Figure 4.3 illustrates a 4-bit MGA representation which requires 4 bits for each individual element of the phenotype. With a 4-bit MGA representation ($c = 4$) the size of the translation table is 2^4 . Figure

4.4 shows a 6-bit MGA representation, where $c = 6$ and the size of the translation table is 2^6 . Finally, Figure 4.5 represents an 8-bit MGA representation where $c = 8$ and the translation table is 2^8 . The granularity of neutrality which is related to the arity of the representation, as is the accessibility of phenotypes from genotypes. The coarseness of the granularity becomes finer as the level of arity increases and visa-versa.

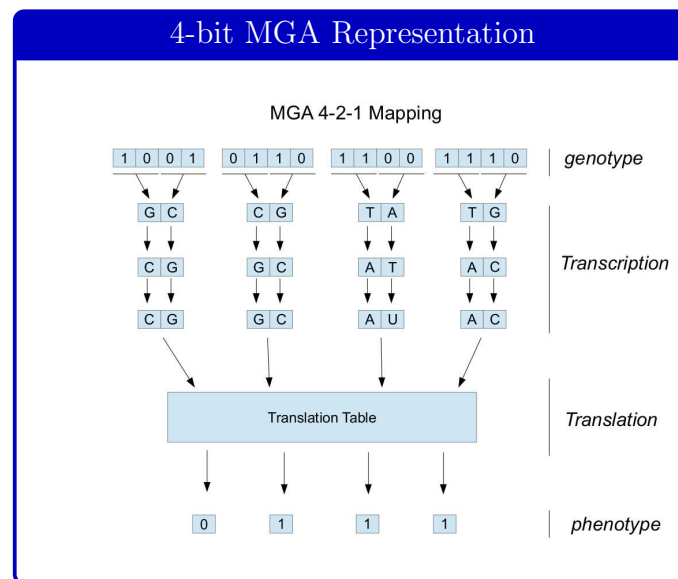


Figure 4.3: 4-Bit MGA Representation Mapping

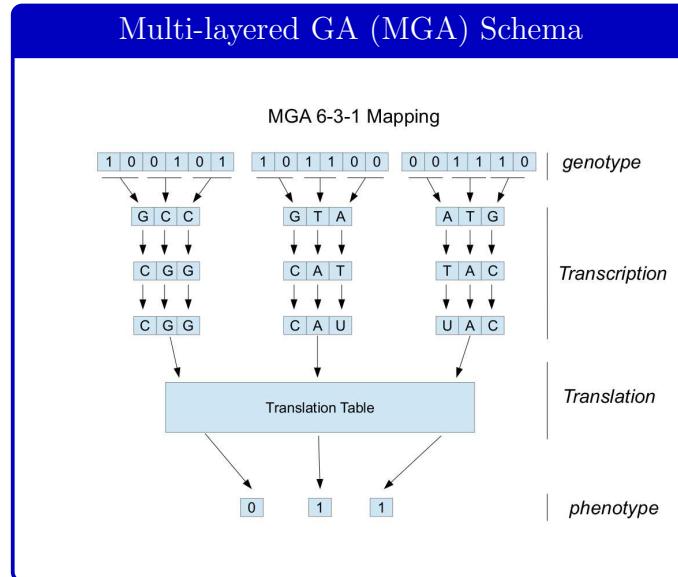


Figure 4.4: 6-Bit MGA Representation Mapping

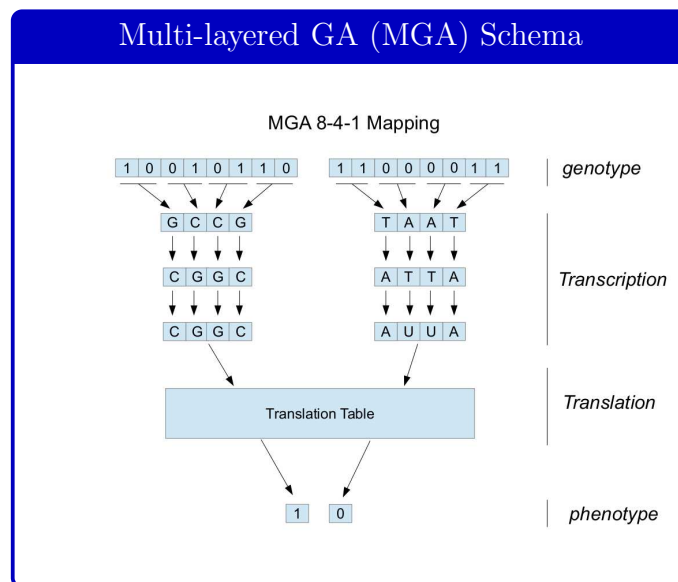


Figure 4.5: 8-Bit MGA Representation Mapping

4.5 MGA Mapping Overview

MGA Mapping Example

Table 4.8 illustrates the mappings from genotype strings which are transcribed and translated into phenes for a 3-bit problem. The mapping process begins with a genome string, which is mapped to the template strand using the template mappings in Table 4.1. This is then converted to create the coding strand and following this the next phase in the transcription stage is to create RNA. This is achieved by using the RNA mappings in Table 4.1, essentially once the DNA coding string is created, each occurrence of *T* is replaced by *U* to create the RNA string. Once we have created the RNA, the transcription stage is now complete.

Genotype-Phenotype Mapping			
Genotype	100000	101100	110000
Transcription Phase:			
DNA Template	GAA	GTA	TAA
DNA Coding	CTT	CAT	ATT
RNA	CUU	CAU	AUU
Translation Phase:			
Phenotype - Translation table mapping	1	1	1
Dynamic Activation Implemented	UU		UU
Phenotype - Following Dynamic Activation	0	1	0

Table 4.8: MGA Genotype-Phenotype Mapping

Translation now begins and the RNA is compared to the values contained in the translation Tables Table 4.3, which check for neighbourhood equivalence and then map to phenes. In this example there are matching pairs of RNA located at predefined locations - *UU* - therefore the values are subject to dynamic activation encodings and are altered accordingly with the nearest phenes having their values changed.

Once the phenes have been created they are then combined together to create a 3-bit phenotype as illustrated in Table 4.8, which shows an 6-bit MGA representation.

Missense Mutation Mapping			
Genotype	10000100	10001101	10001111
DNA Template	GACA	GATC	GATT
DNA Coding	CTGT	CTAG	CTAA
RNA Map	CUGU	CUAG	CUAA
Phenotype	1	0	0
Example 1:			
Missense Mutation	CUG C	CUAG	CUAA
Phenotype	1	0	0
Genotype	10001 101	10001101	10001111
Example 2:			
Missense Mutation	CUGU	CUA A	CUAA
Phenotype	1	0	0
Genotype	10000100	10001 100	10001111
Example 3:			
Missense Mutation	CUGU	CUAG	CUA C
Phenotype	1	0	1
Genotype	10000101	10001100	100000 01
Example 4:			
Missense Mutation	CUGU	CUAG	G UAA
Phenotype	1	0	0
Genotype	10000101	10001100	1000 1111

Table 4.9: MGA Missense Mutation Mapping

Missense Mutation Example

Table 4.9 illustrates, using an 8-bit representation, the impact of missense mutation using a number of examples where a mutation takes place. The table begins by mapping a genotype through the transcription and translation phases and producing the phenotype 100. Example 1 shows the result of missense mutation where the RNA map *CUGU* is mutated to *CUGC* and how this is reflected in the genotype. It should be noted that although the genotype has been altered, the phenotype has remained the same as both RNA map values (*CUGU* and *CUGC*) are found in the

same translation table. In Example 2, the RNA map $CUAG$ is mutated to $CUAA$. The effect of this is shown in the genotype change but the phenotype still remains unchanged as $CUAG$ and $CUAA$ are found in the same translation table. In Example 3, the RNA map $CUAA$ is mutated to $CUAC$, the effect of this is to alter both the genotype and the phenotype as the RNA maps $CUAA$ and $CUAC$ are located in different translation tables. Finally, Example 4 shows the RNA map $CUAA$ being mutated to $GUAA$. This is again reflected back onto the genotype but the phenotype remains unchanged as $CUAA$ and $GUAA$ are both members of the same translation table. With missense mutation whether a mutation is silent or adaptive is dependent on the contents of the translation table. In other words, a missense mutation can be adaptive at a genotypic level, but whether it is silent or adaptive at a phenotypic level is dependent upon the translation table and dynamic activation.

MGA Mapping Summary

To summarise $|\phi_g| = \{0, 1\}^{l_g}$ where l_g is the genotype length. The transcription phase maps $|\phi_g| \rightarrow |\phi_d| \rightarrow |\phi_r|$, where $|\phi_d| = \{A, C, G, T\}^{l_g/2}$ with the following mappings: $00 \rightarrow A$; $01 \rightarrow C$; $10 \rightarrow G$ and $11 \rightarrow T$. A bijective mapping maps $|\phi_d| \rightarrow |\phi_r|$, where $|\phi_r| = \{A, C, G, U\}^{l_g/2}$. U is included for biological plausibility and has no impact on the evolution unless we include operators at this level. Following transcription, the translation phase takes place, mapping the RNA space to a phenotype space $|\phi_p|$, $|\phi_r| \rightarrow |\phi_p|$, where: $|\phi_p| = \{0, 1\}^{l/c}$. The value c represents the cardinality chosen at initialisation to create a translation table. The translation table maps $c/2$ characters to a phenotypic bit, either 0 or 1. The level of redundancy is

determined by c and implies $|\phi_g| > |\phi_p|$ where $c > 1$ and c is even. Missense mutation in nature is carried out at the RNA level. In relation to the MGA, the Missense mutation mapping is as follows: $A \rightarrow C, G$ or U , $C \rightarrow A, G$ or U , $G \rightarrow A, C$ or U and $U \rightarrow A, C$ or G . The variation operators, one-point crossover and single-point mutation occur at the genotype level prior to transcription and missense mutation takes place in the RNA space before translation.

4.6 Chapter Summary

This chapter outlined the Multi-layered GA (MGA) which uses the biological concepts of transcription and translation to implement a modular, non-trivial genotype-phenotype mapping which is designed to introduce neutrality into the representation. The MGA, is designed to allow flexibility to exist in the phenotypic variational topology. To achieve this the size of the alphabet to be created is decided at initialisation and each possible ordering must be assembled in the translation table to represent either a '0' or a '1', thereby guaranteeing that the representation will result in a binary mapping. The multiple mappings and translation tables required for the chosen arity of representation are then created by the MGA. The MGA randomly initialises the binary genome string and this in turn is converted using a series of mappings to convert genes into a combination of amino acids which are examined for neighbourhood equivalence in the translation phase and create a phene. An important point to note is that these neighbourhoods are designed to introduce variation in the population and variability into the phenotypic topology, thereby allowing a phenotype to be encoded by a number of genotypes and that the phenotypic evolution can only be

understood by monitoring the neutral traits. The phenes are then combined to form the phenotype. The primary structure of the representation is designed to increase variation and variability.

The mapping feature of the transcription phase is a form of complementing and appears in many ways, a form of inversion. This is an attempt to, in a basic way, recreate the bridging between the gene and the phene, as in natural RNA which is complementary to the template strand. Therefore, the mappings are included to mimic closely, but still at a basic level, naturally occurring phenomenon. The translation phase on the other hand, takes the output from transcription and maps it to a phene through neighbourhood equivalence with the translation tables generated by the MGA. This allows for neutral traits to be included and are of importance as they allow increased variation and induce phenotypic variability. Finally, the structure of the MGA allows for the introduction of operators at levels other than the genotypic level. The following chapter, Chapter 5 looks at the efficacy of the MGA over a number of different landscapes.

Chapter 5

Examining the Efficacy of a Multi-Layered GA

5.1 Introduction

This chapter examines the efficacy of the MGA using a number of experiments designed to test and analyse the impact of a biologically inspired modular, non-trivial genotype-phenotype map, which introduces neutrality into the representation of a GA (as outlined in Chapter 4). The proposed genotype-phenotype map introduces a more flexible phenotypic structure and a higher degree of phenotypic variability through the presence of neutrality. The multi-layered genotype-phenotype map is tested over deceptive landscapes, changing deceptive landscapes and De Jong's test suite, which have all been used in past research as benchmark tests for GAs. The chapter is laid out as follows: Section 5.2 gives an overview of the efficacy experiments conducted. Section 5.4 describes the Test 1 experiments conducted over static and changing

deceptive landscapes and includes extracts from [86, 87, 88]. Section 5.5 outlines the Test 2 experiments carried out over De Jong's test suite and analyses the results, with extracts found in [89]. Finally, section 5.6 summarises and concludes the chapter.

5.2 Efficacy Experiments Overview

The experiments outlined in this chapter are separated into two tests: Test 1 and Test 2. The experiments carried out under Test 1, outlined in Section 5.4, are conducted over a number of static and dynamic fully deceptive landscapes. The intention is to examine the impact of neutrality, introduced through the modular, non-trivial, many-to-one genotype-phenotype map over static and changing deceptive landscapes, each containing a different degree of problem difficulty. Deceptive landscapes were chosen as they have often been used in past research [70, 73, 198, 138] etc. to test the exploration ability of GAs and overcome the difficulty of being drawn towards a deceptive attractor. The experiments conducted in Test 2 are carried out using De Jong's test suite [43], described in Section 2.6.1 and contain a number of problems designed to include features often found in optimisation problems. De Jong's test is often used in determining the robustness of a GA and has become a standard benchmarking test, used to examine the performance of various GAs. Overall, the aim of the experiments outlined in Section 5.4 and Section 5.5 are designed to illustrate the efficacy of the MGA.

5.3 Statistical Significance

The two sample t-test simply tests whether or not two independent populations have different mean values on some measure. The choice of significance level at which you reject a hypothesis H_0 is arbitrary. The Wilcoxon signed-rank test is a non-parametric statistical hypothesis test which is often used to compare two related samples. It has been used as an alternative the paired t-test when the population cannot be assumed to be normally distributed. A t-test and the Wilcoxon rank sum test were used to test for statistical significance between the results obtained with the SGA and the MGA. In this dissertation, results where $p < 0.05$ are referred to as statistically significant and statistically highly significant if they return $p < 0.001$.

5.4 Test 1 - Deceptive Landscapes

To examine and compare the performance of both a simple GA (SGA) and the MGA, experiments were conducted over a number of fully deceptive landscapes which fall into the class of GA-hard problems [70] and are as follows:

- 3-bit fully deceptive experiments [65]
- 10 3-bit fully deceptive experiments [73]
- 10 3-bit changing landscape fully deceptive experiments [87]

The smallest 3-bit fully deceptive problem is similar to that outlined by Goldberg in [65]. The loosely ordered 10 3-bit problem is as described by Goldberg & Bridges

[73] and incorporates the 3-bit problem into 10 3-Bit deceptive order-three sub functions. By running a GA on a static environment it's *adaptiveness* is difficult to test. However, by changing the environment we have an opportunity to examine the ability of a GA to readjust the search and look for a new global optimum, thereby observing the balance between exploration and exploitation. To investigate the adaptability of the MGA, a changing 10 3-bit problem is used which alters the landscape during the search, allowing the opportunity to examine, the performance of both GAs over a dynamic environment, which exist in many real world problems.

The aim of Test 1 is to examine the efficacy of a modular, non-trivial biologically inspired, multi-layered GA over a number of fully deceptive landscapes. By comparing the performance of a SGA over both static and changing fully deceptive environments, with that of multi-layered mapping GA (MGA) the following hypothesis (H_1) is tested: *Including a biologically inspired, multi-layered, many-to-one genotype-phenotype map into a GA, benefits searching fully deceptive changing landscapes.*

5.4.1 Test 1 Experiment Results

Parameters

The experiments were conducted with a population size of 200, apart from the 3-bit deceptive problem, which because of its limited search space, had a population of 20. The population size of 200 was chosen to ensure adequate exploration of the search space and give both GAs a reasonable chance of success. A crossover rate of 0.7 and a mutation rate of $1/l$, where l is the length of the chromosome were used. The 3-bit experiment was run over 100 generations and results averaged over 10 runs.

The 10 3-bit deceptive problem over 3000 generations and averaged over 10 runs and the changing 10 3-bit deceptive problem over 25000 generations and performance averaged over 10 runs. The results of Test 1 are outlined below.

3-bit Deceptive Problem

To gain an initial understanding of the efficacy of the MGA, the off-line performance (average best fitness) for both the SGA and the MGA were monitored. The first set of experiments were run over a 3-bit fully deceptive problem. Figure 5.1 shows the average best fitness for both the SGA and the MGA, with both variations of the GA locating the global optimum (maximum fitness level of 30) quite quickly as the 3-bit deceptive problem is relatively easy and both algorithms had little difficulty in solving it.

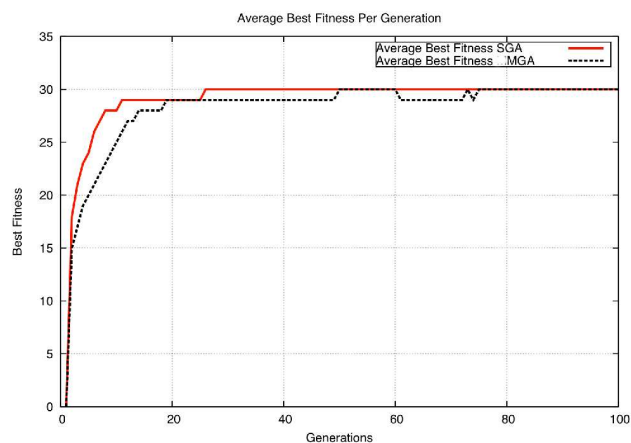


Figure 5.1: Off-line Performance (Average Best Fitness) 3-Bit Deceptive Problem

Overall, when we examine Figure 5.1 it is difficult to argue for the introduction of the layered genotype-phenotype mapping proposed in the MGA given the extra com-

putational overhead associated with implementing the proposed genotype-phenotype map, which seems similar to the findings from Doerr et al. [49] and Collins [34]. So although the MGA exhibits the ability to solve the problem, in the case of the 3-bit deceptive problem, it would appear that this landscape is relatively easy to solve, therefore we need to examine a more difficult landscape to see if there is any benefit associated with the introduction of the neutrality through the proposed layered mapping.

10 3-bit Deceptive Problem

Figure 5.2, illustrates the off-line performance (average best fitness) achieved by both the MGA and SGA per generation, over a 10 3-bit fully deceptive landscape. The optimal fitness for this landscape is a value of 300. The results indicate that the MGA's primary structure allows the search to successfully locate the optimum. However, the SGA struggles with the increased difficulty associated with the loosely ordered problem and is drawn towards a local optimum. These findings would suggest that the benefit of introducing a layered neutral mapping is similar to [62], where the authors argued that neutrality is beneficial when it improves the likelihood of sampling the global optimum.

Figure 5.3, represents this from a different perspective, by outlining the number of sub-functions discovered by both the SGA and MGA, with each sub-function representing a loosely ordered three-bit problem as outlined in Chapter 2. The results indicate that the SGA never manages to locate the entire ten sub-functions, whereas the MGA succeeds in discovering all ten sub-functions, which corresponds to

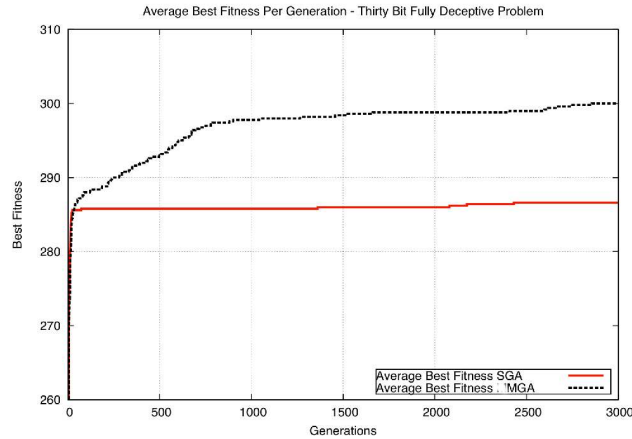


Figure 5.2: Off-line Performance (Average Best Fitness) 10 3-Bit Deceptive Problem

the global optimum. To ensure that the results were statistically significant a t-test analysis was carried with a 95% confidence interval, the result of which indicated that the findings were statistically highly significant, with a p -value score of < 0.001 .

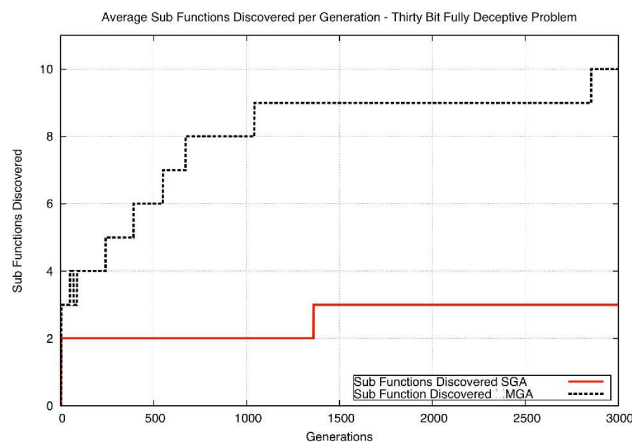


Figure 5.3: Sub-Functions Discovered 10 3-Bit Deceptive Problem

From the results outlined in Figures 5.2 and 5.3, the introduction of a modular, non-trivial genotype-phenotype map appears to be beneficial in searching this particular class of problem. These results concur with Beaudoin [18] which found that by introducing neutrality easier problems became more difficult and more difficult problems became easier. To continue testing the efficacy of the proposed neutral representation a changing landscape is now introduced and examines the adaptive qualities of the MGA and its ability to maintain a balance between exploration and exploitation.

Changing Landscape 10 3-bit Deceptive Problem

As seen from the experiments outlined above, the MGA possesses the capability to solve problems over a relatively difficult deceptive static landscape. To observe the effects of a changing environment the fitness function was altered at a predefined point in the search. This changes the landscape and creates a new global optimum, which in turn examines the adaptive qualities of the algorithms and ensures that an adequate amount of exploration is maintained during the search. In testing the efficacy of the MGA over a changing deceptive landscape, the aim is to increase our understanding of the impact which the multi-layered genotype-phenotype map has on the balance between exploration and exploitation, as the structure of this landscape should give an advantage to the SGA. The reason for this is because the SGA gets drawn towards the deceptive attractor in the first stage of the search and when the landscape changes, the deceptive attractor becomes the new global optimum, therefore the SGA begins the search closer to this point.

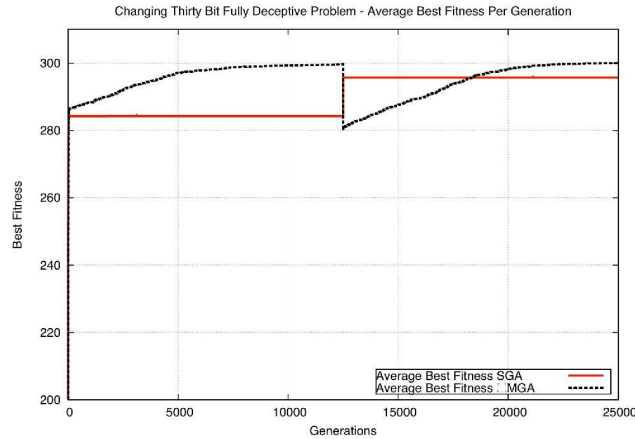


Figure 5.4: Off-line Performance 10 3-Bit Changing Landscape

Figure 5.4 shows the off-line performance for both algorithms. Once the search reaches the half way point (generation 12,500), the fitness function changes and the landscape shifts where the global optimum prior to the changing of fitness functions, switches and becomes the deceptive attractor of the new objective function. At this point the SGA is closer to the new global optimum, while the MGA is located directly on the deceptive attractor. As the search continues the SGA remains trapped on the local optimum and fails to locate the optimum (maximum fitness value of 300). The MGA however, succeeds in locating the global optimum before the landscape changes and adapts to escape the local optimum allowing it to locate the new global optimum after the landscape changes.

Figure 5.5, illustrates the number of sub-functions discovered and shows that in the first half of the search, the SGA is drawn away from the global optimum and converges on the deceptive attractor, solving only, on average, two sub-functions. The

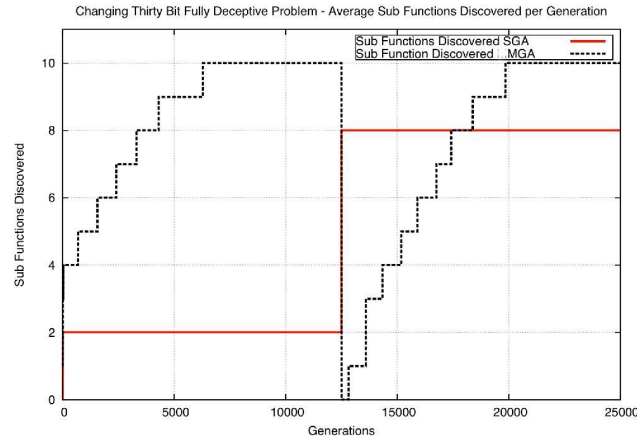


Figure 5.5: Sub-functions Discovered 10 3-Bit Changing Landscape

MGA on the other hand, continues probing the landscape and discovers the global optimum, solving all ten sub-functions. When the landscape changes, the SGA is closer to the global optimum as its performance was poor prior to the change in fitness values, but remains trapped and fails to locate all ten sub-functions. While the MGA continues to search after the landscape has changed and succeeds in optimising the problem and solving all ten sub-functions. The results indicate that the layered representation included in the MGA improves the adaptability of the algorithm over changing landscapes, maintaining a balance between exploration and exploitation. The findings are similar to those of Ebner et al. [52, 51] who found neutrality advantageous over changing landscapes.

5.5 Test 2 - De Jong's Test Suite

To continue examining the efficacy of the MGA, De Jong's [43] test suite was chosen as it is often used as a benchmark for GA performance. The results outlined in the experiments below, illustrate both the average performance, that is the measure of the average fitness of all members of the population (on-line performance), and the average best performance, that is the average maximum fitness of the population (off-line performance). The aim of Test 2 is to test the hypothesis (H_2) *The inclusion of a biologically inspired, modular, fixed, non-trivial, multi-layered genotype-phenotype map into a GA, can perform as robustly as a standard GA in optimising De Jong's test suite.*

5.5.1 Test 2 Experiment Results

Parameters

The parameters used for the Test 2 experiments are as follows; a crossover rate of 0.70, a mutation rate of $1/l$ where l is the length of the genotype and a population size of 200. The number of generations varied for each set of experiments and are as follows: the Sphere Model ran for 100 generations, Rosenbrock's function for 2000 generations, the Step function for 500 generations, the Quadratic function with noise for 500 generations and finally Shekel's function for 500 generations. The results of the experiments carried out are each averaged over 50 runs and are discussed below.

The Sphere Model (f_1)

The results of the Sphere model (outlined in 2.6.1) experiments are shown in Figure 5.6 and illustrate the findings for both the SGA and the MGA, showing the off-line performance and on-line performance for each generation. The problem is three dimensional and the optimum is achieved where $x_1 = 0$, $x_2 = 0$ and $x_3 = 0$. In the experiments conducted, both the SGA and the MGA easily locate the global optimum.



Figure 5.6: The Sphere Model - SGA & MGA.

Both the off-line and the on-line performance for the SGA are similar, indicating that with the SGA, as the population converges towards the fittest individual, the average is also converging towards the fittest level. The MGA's off-line performance is marginally better than that of the SGA and although not visible in the figures, the earliest, on average, location of the optimum for the SGA is during generation 83, while the global optimum for the MGA is located much sooner, on average during generation 25. However, the MGA's on-line performance falls short of the SGAs, due

to the level of neutrality present in the MGA's representation, indicating that the inclusion of neutrality in the representation has had an impact on variation within the population. Overall, the differences between the SGA and MGA on this problem are negligible. A Wilcoxon rank sum test was used to test for statistical significance. The statistical tests carried out on the on-line performances of both the SGA and MGA were statistically highly significant, with a p -value $< 2.2e^{-16}$. Similarly, the statistical tests carried out on the off-line performances indicated that the results were also shown to be statistically highly significant (p -value $< 2.462e^{-15}$).

Rosenbrock's Function (f_2)

The second set of experiments were carried out over Rosenbrock's function (outlined in 2.6.1), with the intention of testing the performance of the algorithms in dealing with the direction of the search repeatedly changing. The global optimum for this two dimensional problems is where $x_1 = 1$ and $x_2 = 1$. Figure 5.7 illustrates the performance of both GAs. The global optimum is located by the SGA, on average over the 50 runs, at generation 249 as the ridge is relatively easy to locate. The difficulty associated with this function lies in its ability to converge. The on-line performance of the population indicates the level of convergence. As was the case with the Sphere Model experiments, the off-line and on-line performance converges as the search progresses. Looking at the off-line performance of the MGA, it has located the global optimum quite quickly (on average during generation 99). This compares favourably with the SGA, indicating far fewer function evaluations for the MGA in locating the optimum.

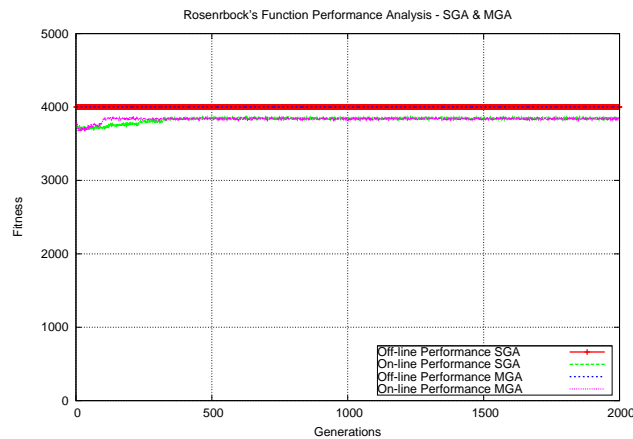


Figure 5.7: Rosenbrock's Saddle - SGA & MGA

Although the off-line performances are close, it is interesting to note that the on-line performances appear quite similar for both the SGA and the MGA. One possible reason for this is that due to the nature of the search space, convergence is not as easy to obtain as was the case with the Sphere function, and both the SGA and the MGA experience similar difficulties. Overall, the off-line and on-line performance for both the SGA and the MGA are relatively similar over this search space. Statistical results using a Wilcoxon rank sum test, show that both the on-line and off-line performances of both GAs were statistically highly significant with a p -value $< 2.2e^{-16}$.

Step Function (f_3)

The Step function, discussed in 2.6.1, highlights an algorithm's ability not to get trapped on a plateau of local optima. From Figure 5.8 we see that the SGA performs well and discovers the global optimum (with a fitness level of 30) on average at generation 203, indicating the SGA's off-line performance over the landscape.

Also, both the off-line and on-line performances for the SGA are almost identical




Figure 5.8: Step Function - SGA & MGA

at an early stage in the search. However, the MGA discovers the global optimum, on average, during generation 4, which is a significant improvement over the off-line performance of the SGA. It would appear that the neutrality present in the MGA representation is beneficial as it appears to improve the probability of sampling the global optimum [63], by increasing exploration while still maintaining a balance with exploitation. With regard to the on-line performance, there is little difference between both algorithms over this problem domain. Again, the Wilcoxon rank sum test illustrated that the SGA and MGA on-line and off-line results were statistically highly significant ($p\text{-value} < 2.2e^{-16}$).

Quadratic Function with Noise (f_4)

The quadratic function with noise (outlined in 2.6.1) was designed to evaluate the performance of an algorithm over a landscape which continuously changes due to the presence of noise. The results of the experiments are shown in Figure 5.9 and although difficult to identify, in the figure, the SGA's best off-line performance occurs

on average at generation 492, illustrating the difficulty associated with noise in the search space. The on-line performance tends to vary per generation, again due to the presence of noise.




Figure 5.9: Quadratic Function with Noise - SGA & MGA

The performances both off-line and on-line for the MGA look quite similar to that of the SGA, with the best off-line performance for the MGA occurring on average at generation 440. However, the MGA recorded a higher off-line performance earlier in the search, which may indicate that in the presence of noise, there may be a benefit associated with the inclusion of the type of neutral representation introduced by the MGA, as it improves the likelihood of reaching the global optimum [62]. It may also indicate that through the dampening of the destructive effects of variation operators, similar to [191, 207], the MGA representation may be beneficial in a noisy environment. The statistical results, using a Wilcoxon rank sum test showed that the off-line results of the SGA and MGA were statistically highly significant with a p -value $< 2.2e^{-16}$. However, the on-line results were not significant, with a p -value < 0.2034 .

Shekel's Foxhole's Function (f_5)

The final set of experiments compare both algorithms over Shekel's Foxhole's function (outlined in 2.6.1), with the results illustrated in Figure 5.10. Both the SGA and the MGA solved the problem; the SGA locates the global optimum, on average, during generation 266, while the MGA, on the other hand, locates the global optimum on average at approximately generation 50.




Figure 5.10: Shekel's Foxholes - SGA & MGA

The results indicate that there is a significant improvement in performance over the multi-modal landscape of Shekel's Foxholes when an element of neutrality is incorporated into the representation. It appears that the MGA possesses the ability to avoid getting trapped in local optima as it maintains a balance between exploration and exploitation. One possible reason for this could be that the neutrality introduced through the multi-layered mapping, reduces the impact of operators such as mutation, similar to [191, 207] as only adaptive mutation alters fitness, silent mutations maintain diversity [205], assisting exploration and proves beneficial in avoiding getting stuck

in a local optima as the search progresses. By partially insulating against the effect of the operators the translation table assists in maintaining within the population, valuable building blocks, allowing escape routes from local traps to be found [54]. Overall, the MGA has both a better off-line and on-line performance over the SGA for the type of search space produced by Shekel's function. The results of both GAs were statistically highly significant as a Wilcoxon rank sum test returned a p -value $< 2.2e^{-16}$, for the on-line and off-line performances of the SGA and MGA.

5.6 Chapter Summary

The experiments outlined in this chapter were designed to test the efficacy of the MGA whose genotype-phenotype mapping introduces neutrality into the representation, over deceptive landscapes, changing deceptive landscapes and De Jong's test suite. The Test 1 experiments were designed to test the hypothesis H_1 *Including a biologically inspired, multi-layered, many-to-one genotype-phenotype map into a GA, benefits searching fully deceptive changing landscapes.* The experiments moved from a basic 3-bit fully deceptive problem to a loosely ordered 10 3-bit fully deceptive problem and finally, to a 10 3-bit fully deceptive changing landscape problem, with results indicating that the proposed multi-layered genotype-phenotype map is effective in solving the deceptive problems presented. The results found that there was little benefit, if any, of introducing neutrality for the 3-bit deceptive problem, mainly due to the relative ease of the problem, as found in Collins [34]. These results also correspond with Beaudoin et al. [18], who found that as the level of difficulty increased so too did the benefit of including neutrality and Doerr et al. [49], who

found that neutrality was possibly beneficial over more difficult deceptive landscapes with multiple local optima as run-time reduced from exponential to polynomial. Furthermore, the MGA representation appears to be structured in a way that improves the likelihood of sampling the optimum [62], combined with the dampening of the destructive effects of mutation [191, 207]. Therefore, as the problems moved from static deceptive problems; to more difficult deceptive problems; to dynamic deceptive landscapes, the benefits associated with the inclusion of the mapping continued to increase. The Test 1 results indicate that the benefits associated with the multi-layered the stated hypothesis (H_1) is supported.

To continue examining the efficacy of the MGA, Test 2 experiments were designed to test the following hypothesis H_2 *The inclusion of a biologically inspired, modular, fixed, non-trivial, multi-layered genotype-phenotype map into a GA, can perform as robustly as a standard GA in optimising De Jong's test suite.* The experiments conducted for Test 2 show that the MGA performed as robustly as the SGA. However, for the characteristics present by the Sphere function, the Rosenbrock function, the Step function and the Quadratic function, the benefit of neutrality is not apparent at first sight and for many it is negligible, but this was not the case for the Shekel's Foxholes experiments, where the introduction of neutrality into the genotype-phenotype mapping allowed the MGA to outperform the GA. A possible reason for this is that the neutrality introduced through the multi-layered mapping, reduces the impact of operators such as mutation and crossover and slows genetic drift, as found in [191, 207]. It also seems, from the results presented, that the level of problem difficulty is a significant factor in deciding whether or not the inclusion of the MGA's neutral rep-

resentation would be beneficial, which concurs with [18]. The results of Test 2 indicate the hypothesis proposed (H_2) is supported.

In summary, this chapter has illustrated and established the efficacy of the multi-layered GA (MGA) and that the MGA outperformed the SGA over the more difficult problems and illustrates its advantages over a deceptive changing landscape. Combined with the dampening of the effects on mutation as only adaptive mutation alters fitness, with silent mutations maintaining diversity [205], which assists exploration and appears beneficial as the search progresses. This change in variation is critical and as the proposed mapping was designed to maintain variation and increase variability, a question remains as to what their impact is on the search process. Maintaining variation and inducing variability are also possible reasons as to why the proposed mapping proved beneficial over the class of problems tested, but to isolate their impact, further experimentation is required. To begin to understand the influence of the introducing a neutral mapping into the representation of a GA on variation and its impact on influencing variability, an examination of diversity, both at a genotypic and phenotypic level is required. With this in mind the focus of the next chapter, Chapter 6, is on the impact of the proposed mapping on population variation.

Chapter 6

Examining Population Variation in a Multi-layered GA.

6.1 Introduction

Results from Chapter 5 established the efficacy of the MGA and indicated that the presence of neutrality within the primary structure of a GA was beneficial in searching more difficult and changing landscapes. However, as the proposed many-to-one neutral mapping was designed to maintain variation and increase variability, these characteristics need to be examined to establish reasons as to why the MGA performed well over the class of problems chosen. This chapter examines the variation associated with both the SGA and the MGA. Variation can be described as the difference between individuals in a population and in this dissertation, is interpreted as the level of diversity present in the population. The chapter, using a number of experiments, examines the impact on variation caused by the inclusion of the pro-

posed neutrality within the primary structure of a GA. The chapter focuses on the maintenance of diversity within the population, problem difficulty and inducing phenotypic variability. Extracts of this chapter are taken from [88, 90] and the chapter is laid out as follows: Section 6.2 gives an overview of the experiments, Section 6.3 describes Test 3 and examines the impact on variation of altering the level of problem difficulty on population diversity over a number of fully deceptive landscapes. Section 6.4 outlines Test 4 and includes a number of well known problems, where the degree of dimensionality is altered to increase the level of difficulty associated with each problem and examines the effects on variation. Finally, Section 6.5 summarise and concludes the chapter.

6.2 Variation Experiments Overview

The experiments conducted in this chapter are separated into two tests, Test 3 and Test 4. Test 3 accesses the levels of variation present in the population and its impact on variability, with experiments carried out over deceptive and changing deceptive landscapes, that is the 3-bit, 10 3-Bit and changing 10 3-Bit deceptive landscapes. These experiments can be viewed as an extension to those carried out in Chapter 5, but focus on variation rather than efficacy. Test 4 looks at the relationship between variation and problem difficulty and examines the impact on variation of changes in the degree of problem difficulty over a number of well understood problems. The problems chosen were the Sphere function [43] as outlined in Section 2.6.1 and a changing Sphere function [94], described in Section 2.6.2. The aim of the chapter is to gain a better understanding of how the proposed genotype-phenotype mapping,

included in the MGA, impacts on variation in the population, and whether this assists in searching the search space of the selected landscapes.

6.3 Test 3 - Inducing Phenotypic Variability

Test 3 examines a number of fully deceptive, noisy and dynamic problems. The aim of Test 3 is to examine the impact on population diversity caused by the proposed MGA representation. Put another way, how does the inclusion of neutrality, resulting from the MGA's genotype-phenotype map, impact on population variation. Test 3 is designed to test the following hypothesis (H_3): *Diversity, implicitly maintained by a many-to-one, genotype-phenotype mapping of a GA, implemented by an interpretation of the biological processes of Transcription and Translation, is beneficial in searching noisy and dynamic landscapes.* The experiments used for Test 3 are as follows:

- 3-bit fully deceptive experiments [65]
- 10 3-bit fully deceptive experiments [73]
- 10 3-bit changing landscape fully deceptive experiments [87]

For each group of experiments, diversity is calculated at both a genotypic and phenotypic level using pair-wise hamming distance in order to gain an understanding to the level of variation present in the topology. For comparison purposes, the data was normalised on the hamming distance measurements and uses a scale of 0 to 1, with 0 indicating convergence and 1 indicating maximum diversity.

6.3.1 Test 3 Experiment Results

Parameters

The experiments were conducted with a population of 200, apart from the 3-bit deceptive problem, which because of its limited search space, had a population size of 20. The population size of 200 was chosen to ensure adequate exploration of the search space and give both GAs a reasonable chance of success. A crossover rate of 0.7 and a mutation rate of $1/l$, where l is the length of the chromosome were used. The 3-bit experiment was run over 100 generations and averaged over 10 runs. The 10 3-bit deceptive problem over 3000 generations and averaged over 10 runs and the changing 10 3-bit deceptive problem over 25000 generations and averaged over 10 runs. The results of Test 3 are outlined below.

3-bit Deceptive Problem

To develop an understanding of what effect the MGA genotype-phenotype map would have on variation, the diversity of both the SGA and MGA are monitored. The results of the 3-bit deceptive problem [65] (outlined and discussed 5.4.1) indicate that for both the SGA and MGA the landscape was relatively easy for both algorithms. Figure 6.1 illustrates the level of genotypic and phenotypic diversity in the population for both the MGA and a SGA. The SGA, with a one-to-one genotypic-phenotypic representation, converges very quickly and diversity, both genotypic and phenotypic diversity are the same, apart from the efforts of the mutation operator. In relation to variation, once convergence has occurred the result is a lack of both genotypic and phenotypic variation. The MGA, through the modular, non-trivial mapping,

maintains a level of diversity within the primary structure resulting in a maintained level of variation. Figure 6.1 illustrates the presence of variation both at a genotypic level and to a lesser extent, at a phenotypic level. These results in relation to the maintenance of diversity, are similar to results reported by Yu and Miller [205] and Banzhaf [10].

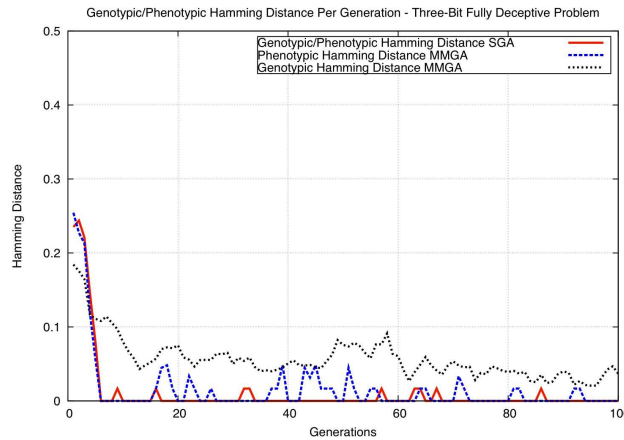


Figure 6.1: Genotype/Phenotype Diversity 3-Bit Deceptive Problem

A paired t-test was used to determine if the results of the hamming distance experiments between the SGA and the MGA. The results of the paired t-test between the normalised SGA phenotype data and the normalised MGA phenotype data indicated that the true difference in means is not equal to 0 and the p -value was 0.001656, indicating the results were statistically significant ($p > 0.001$). The result of the paired t-test carried out on the SGA genotype and the MGA genotype returned a p -value $< 2.2e^{-16}$ indicating that the results were statistically highly significant ($p < 0.001$). For each experiment, the results of the t-test analysis with a 95% confidence interval,

indicates that the result of the 3-bit deceptive problem hamming distance experiments were highly statistically significant with $p < 0.001$ for the SGA genotype and the MGA genotype tests, indicating that both genotypes differ statistically. These results were not as significant $p > 0.001$ (but still significant) for the SGA phenotype and the MGA phenotype tests, indicating that the SGA phenotype and the MGA phenotype are similar, apart from the impact of mutation. Overall the landscape proved very easy for both algorithms, with the SGA performing better than the MGA, even though the MGA maintained more variation. The conclusion here, which is similar to that reported by Beaudoin et al. [18], is that over easier landscapes the benefit of including neutrality is diminished. The next set of experiments, examine the impact of increasing the level of problem difficulty on population variation.

10 3-bit Deceptive Problem

As the efficacy experiments indicated that problem difficulty was a relevant feature, the 10 3-bit fully deceptive problem [73] increases the level of problem difficulty and results outlined in 5.4.1 indicate that the landscape was misleading for the SGA, while the MGA optimises the fully deceptive topology. Figure 6.2, highlights both the genotypic and phenotypic diversity associated with both the SGA and MGA. Again, because of the one-to-one representation and genotype-phenotype map present in the SGA, both the genotypic and phenotypic diversity are the same. The SGA's genotypic and phenotypic diversity converges early on, as it is drawn towards the deceptive attractor, indicating a lack of variation within the population. The MGA's performance improves on that of the SGA, as a level of diversity, both phenotypic and

to a greater extent genotypic, is maintained throughout the search. Variation is also greater when compared to the 3-bit problem, as the variational topology changes. The ability to maintain, implicitly, a level of diversity through the genotype-phenotype map increases the level of variation. This increase in variation appears to assist in searching the more difficult landscape, as the MGA exhibits the ability to repeatedly locate the optimum solution where the SGA fails, indicated by the results of the 10 3-bit deceptive experiments. One reason for this is that the layered mapping, while maintaining variation through a slowing of genetic drift caused by neutrality, induces variability, allowing paths over neutral networks to be discovered, thereby allowing the escape from the deceptive attractor, in a similar manner to findings reported by Elgin [54].

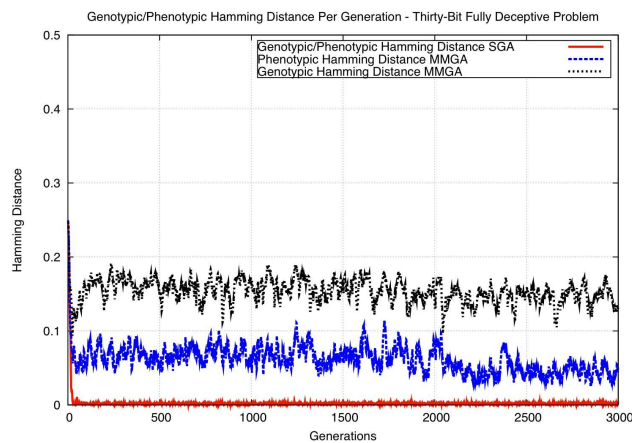


Figure 6.2: Genotype/Phenotype Diversity 10 3-Bit Deceptive Problem

Changing Landscape 10 3-bit Deceptive Problem

The results from the efficacy experiments indicated, the adaptability of the MGA showed significant advantage over the more difficult deceptive changing landscapes. The following experiment, examines variation in order to seek an explanation for the improved performance associated with the MGA. Over the changing 10 3-bit deceptive landscape, results discussed in Section 5.4.1, show that the MGA successfully managed to locate the global fitness, where the SGA struggled. Testing the MGA over a changing deceptive landscape aims to further increase our understanding of the impact the multi-layered genotype-phenotype map has on variation, genetic drift and variability. The structure of this landscape should be advantageous to the SGA, as the SGA gets drawn towards the deceptive attractor in the first stage of the search and at the point at which the landscape changes, the deceptive attractor becomes the new global optimum. Because of this the SGA should begin the search closer to the optimum.

Figure 6.3 indicates that both the SGA's genotypic and phenotypic diversity disappears early on as the population converges. However, with the MGA, we have increased variation brought about by the introduction of neutrality through the genotype-phenotype map which implicitly maintains a level of diversity in the population. A t-test analysis with a 95% confidence interval was used and the results of the hamming distance experiments were shown to be statistically significant. The results suggest that by implicitly maintaining diversity within the population through the multi-layered genotype-phenotype mapping, the MGA increases variation and as reported by Ebner et al. [51, 52], the maintaining of diversity can be advantageous

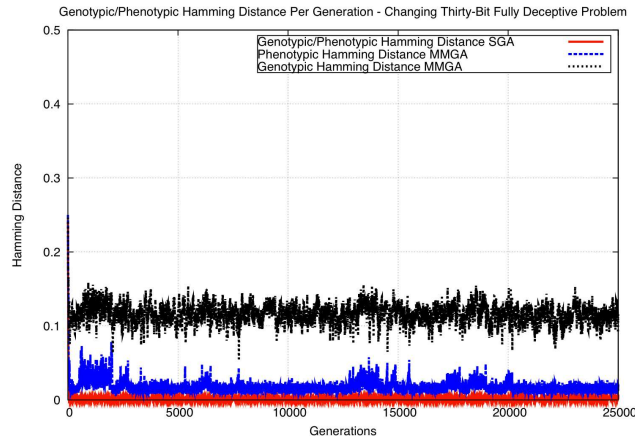


Figure 6.3: Genotype/Phenotype Diversity 10 3-Bit Changing Landscape

in a changing environment.

6.4 Test 4 - Altering Dimensionality

Having examined the impact on variation over static and dynamic fully deceptive landscapes. The aim of Test 4 is to vary the degree of problem difficulty by altering the dimensionality for each problem and measure its impact on population variation. The experiments used in Test 4 were created to test the following hypotheses:

- (H_4) *Introducing neutrality into the representation of a GA, through a modular, fixed non-trivial mapping, is beneficial in searching static landscapes with varying levels of difficulty.*
- (H_5) *Including a modular, fixed non-trivial genotype-phenotype map which introduces a more flexible phenotypic structure and a higher degree of phenotypic*

variation through the use of neutrality, benefits the optimisation of solutions over dynamic landscape problems.

The experiments to test H_4 are carried out over the Sphere Model [43] and the experiments to test H_5 are run over the changing Sphere Model [94].

6.4.1 Test 4 Experiment Results

Parameters

The experiments were carried out with a population of 400 and the parameter values used the experiments for both of the GAs are as follows; crossover rate $P_c = 0.7$, mutation rate $P_m = 1/l$, where l is the length of the chromosome. The Sphere Model experiments ran for 2000 generations and the changing environment experiments ran for 4000 generations.

Sphere Model Static Environment

The Sphere Model [43] outlined in 2.6.1, is relatively easy to optimise as it is continuous, convex and unimodal. This function is normally used to measure the efficiency of a particular algorithm. The nature of the problem allows the ability to altered the number of dimensions (n) associated with the Sphere Model, in order to vary the level of difficulty and to examine the impact of dimensionality on the algorithms. Experiments on the Sphere Model were carried out where $n = 3$, $n = 15$ and $n = 30$.

Sphere Model - Dimensions $n = 3$

Results of the first experiment of Test 4 carried out on the Sphere Model are outlined in Table 6.1 and shows the percentage of times the global optimum is located by the SGA and MGA, along with the average number of runs required to locate the global optimum. The results are averaged over the number of successful runs where the number of dimensions $n = 3$ and illustrate the impact on the performance of the algorithms over the Sphere Model function.

Sphere Model Results		
Number of Dimensions $n = 3$		
GA Description	SGA	MGA
Optimum Located	100%	100%
Avg. No. Generations Required	5	22

Table 6.1: Sphere Model Experiment $n = 3$




Figure 6.4: Sphere Model Diversity $n = 3$ - SGA & MGA


When the number of dimensions n is set to 3 the problem is extremely easy for the SGA and the MGA, with both achieving 100% success in locating the global optimum. The SGA locates the optimum on average after only 5 generations and the MGA takes on average 22 generations (see Table 6.1). One possible reason for this is that an adequate level of diversity in the population exists early in the search and that due to the relative ease of the problem, the SGA on average, locates the global optimum before the MGA. Results reported by Doerr et al. [49] carried out over a unimodal landscape on the effects of incorporating neutrality, suggest that there was no significant advantage of using neutrality. The level of diversity in the population for both the SGA and the MGA is shown in Figure 6.4. It should be noted that the SGA's level of genotypic and phenotypic diversity are one and the same, due to the one-to-one genotype-phenotype mapping, whereas the genotype diversity and the phenotype diversity of the MGA differ because of the nature of the neutral genotype-phenotype mapping. Another consideration is that the introduction of neutrality comes at a cost, as there is extra computational overhead required due to the increase in the search space. Which means that any additional benefit may be outweighed by the additional overhead [63], this is particularly relevant for less difficult landscapes.

Sphere Model - Dimensions $n = 15$

When the number of dimensions are increased to 15, the level of difficulty increases for both GAs. Looking at Table 6.2, the level of difficulty can be seen in the average number of generations taken to locate the global optimum. In this experiment both

the GAs obtained a 100% success rate. However, the SGA needs an average of 1014 generations, while the MGA needs an average of 359 generations, indicating that the increase in dimensionality has increased the level of difficulty for both GAs, but the effect is more pronounced for the SGA. This may be due to the lack of diversity within the population as illustrated in Figure 6.5, which impacts on variation and variability.

Sphere Model Results		
Number of Dimensions $n = 15$		
GA Description	SGA	MGA
Optimum Located	100%	100%
Avg. No. Generations Required	1014	359

Table 6.2: Sphere Model Experiment $n = 15$ Figure 6.5: Sphere Model Population Diversity $n = 15$ - SGA & MGA

Looking at Figure 6.5, the population for the SGA converges quite early in the search and this may account for the average number of generations required to locate

the global optimum. However, as the level of difficulty increases through an increase in the number of dimensions, the level of variation within the MGA's genotypic population increases, which appears beneficial in searching the given space. As Wagner [191] suggests, neutrality assisted in sheltering a system from the impact of mutation and argued that new adaptations, provided by neutrality, can assist the search. The results also indicate that as the level of problem difficulty increases, the benefit of including neutrality also increases, which is similar to findings reported by Beaudoin et al. [18]. By dynamically maintaining variation the MGA offers the potential to induce variability, thereby optimising the problem in fewer function evaluations.

Sphere Model - Dimensions $n = 30$

In the final Sphere Model experiment, the number of dimensions were increased to 30. The effect of this can be seen in Table 6.3 and Figure 6.6. This increase in the number of dimensions has had a significant impact on both of the GAs, but again the SGA's performance shows the largest drop off in performance. The SGA only succeeded in locating the global optimum 80% of the time, while the MGA continues with a 100% success rate.

The other interesting result here can be found in the average number of generations required to locate the optimum, with the SGA requiring, on average, 1681 generations (with the average based on those generations that located the optimum), while the MGA needed, on average, 1007 generations. This is a significant improvement in performance on the part of the MGA when compared to that of the SGA. It would appear that as outlined by Ebner et al. [51, 52], higher levels of mutation could be

sustained by having neutral networks present, which assists in searching more difficult landscapes.

F_1 Sphere Model Results		
Number of Dimensions $n = 30$		
GA Description	<i>SGA</i>	<i>MGA</i>
Optimum Located	80%	100%
Avg. No. Generations Required	1681	1007

Table 6.3: Sphere Model Experiment $n = 30$




Figure 6.6: Sphere Model Population Diversity $n = 30$ - SGA & MGA

The results reported indicate that searching a landscape where the dimensionality has increased, may be assisted by maintaining an element of variation in the population, offering the possibility to induce variability. Figure 6.6 shows an implicit level of genotypic diversity being maintained in the population of the MGA. This may indicate that the improvement in performance is due to useful variation being maintained in the population of the MGA [205], inducing variability. The figure illus-

trates that the SGA's level of Genotype/Phenotype diversity decreases quickly, while for the MGA, although the level of phenotypic diversity diminishes relatively quickly, it always appears marginally higher than that of the SGA as genetic drift associated with the MGA population slows [54], due to the amount of variation maintained by the MGA as a result of the many-to-one genotype-phenotype mapping.

Sphere Model Changing Environment Results

The changing Sphere environment [94] and outlined in Section 2.6.2, has at its heart the Spear model as outlined in [43]. However, the idea is to allow the GAs to search the landscape defined by the sphere model and to then change the function values after 1500 generations, so that the landscape also changes. The aim of this set of experiments is to examine how both the SGA and MGA adapt in a changing landscape environment. The experiments on the Changing Sphere Model were carried out where $n = 3$, $n = 15$ and $n = 30$.

Changing Sphere Model - Dimensions $n = 3$

The results of the first changing Sphere Model experiment are shown in Table 6.4, which shows the percentage of success after the objective function changes. As the dimension level n is set to 3, the problem landscape is relatively easy for both the SGA and the MGA, with both algorithms succeeding 100% of the time. However, the MGA discovers the global optimum in the changing landscape in an average of 1543 generations, which is an improvement in performance over the SGA, which took an average of 2522 generations, indicating that the neutral representation may be beneficial over a changing landscape, as reported by Ebner [51, 52].

Sphere Model Changing Environment		
Number of Dimensions $n = 3$		
GA Description	<i>SGA</i>	<i>MGA</i>
Optimum Located	100%	100%
Avg. No. Generations Required	2522	1543

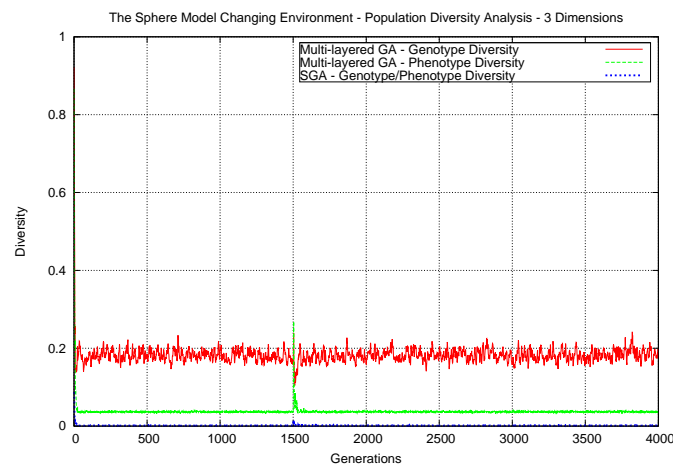
Table 6.4: Sphere Model Changing Environment Experiments Dimensions $n = 3$ Figure 6.7: Sphere Model Changing Environment Diversity $n = 3$ - SGA & MGA

Figure 6.7 illustrates the population diversity for both the SGA and the MGA where $n = 3$. The diversity is similar to that of the first Sphere Model experiment where $n = 3$, as the problem is relatively easy for both algorithms, both before and after the landscape changes. However, due to the level of problem difficulty this change hardly registers within the population variation remaining relatively constant.

Changing Sphere Model Dimensions $n = 15$

By changing the number of dimensions to $n = 15$, Table 6.5 indicates that the SGA is finding it difficult to cope with the changing environment when problem difficulty increases due to the increase in the level of dimensionality. The SGA only manages to succeed in locating the new global optimum 20% of the time after the landscape changes and when it did locate it the average number of generations required was 2939. The MGA, on the other hand, was able to locate the global optimum 100% of the time and the average number of generations required was 2182. Again as with previous experiments, the presence of neutrality appears to make easier problems more difficult and harder problems easier [18]. Also the presence of neutrality slows convergence [36], which appears to benefit more difficult changing landscapes.

Sphere Model Changing Environment		
Number of Dimensions $n = 15$		
GA Description	SGA	MGA
Optimum Located	20%	100%
Avg. No. Generations Required	2939	2182

Table 6.5: Sphere Model Changing Environment Experiments Dimensions $n = 15$

The population variation as, illustrated in Figure 6.8, shows the SGA converging early on in the search as the diversity disappears from the population. The MGA population continues to maintain a level of diversity and the genotypic level and to a lesser extent at the phenotypic level. The level of diversity contained in the MGA

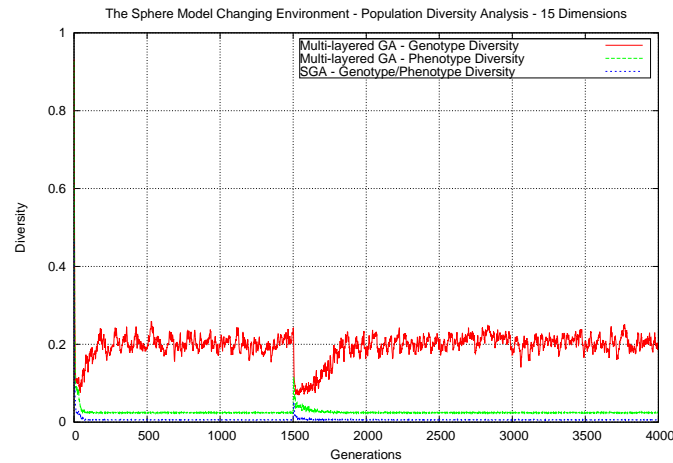


Figure 6.8: Sphere Model Changing Environment Diversity $n = 15$ - SGA & MGA

population begins to increase at the genotypic level, when the search commences and then continues to dynamically maintain diversity within the population. Once the landscape changes, the level of diversity drops and begins increasing again. It then maintains a degree of diversity within the population. One reason for this is that the genotype-phenotype mapping allows for a many-to-one relationship between the genotype and the phenotype, thereby increasing the variation and inducing variability [36], affording the MGA the ability to maintain a balance between exploration and exploitation and adapt in a changing environment.

Changing Sphere Model - Dimensions $n = 30$

The final experiment in the changing environment experiments, saw the level of dimensionality n increasing from 15 to 30. With this level of difficulty the SGA was unable to locate the global optimum after the landscape changed, illustrated by 0% in Table 6.6. The MGA, however, did locate the new global optimum 100% of the

time, with an average number of generations of 3189 required (see Table 6.6).

Sphere Model Changing Environment		
Number of Dimensions $n = 30$		
GA Description	SGA	MGA
Optimum Located	0%	100%
Avg. No. Generations Required	N/A	3557

Table 6.6: Sphere Model Changing Environment Experiments Dimensions $n = 30$




Figure 6.9: Sphere Model Changing Environment Diversity $n = 30$ - SGA & MGA

When we examine the population diversity shown in Figure 6.9 we can see a similar pattern to the previous experiments, in that the SGA population variation disappears in the search and the MGA, through its multi-layered mapping, dynamically maintains a level of variation within the population. One other interesting feature associated with the dynamic nature of diversity maintenance within the MGA population, is that the normalised levels of diversity tend to increase as the level of difficulty

increases. Comparing Figures 6.7, 6.8 and 6.9, it appears that the MGA's level of diversity is increasing in response to the level of difficulty. This dynamic adjustment of diversity would appear a useful characteristic in attempting a search of more difficult dynamic environments, giving the MGA a robustness similar to that found by Wilke [201]. One possible reason for this is that as neutrality dampens the influence of variation operators, genetic drift slows and variation is maintained within the population. These findings exhibit similarities to previous findings [18, 52, 51, 205, 206], in relation to the benefits of neutrality over more difficult and changing landscapes.

6.5 Chapter Summary

The experiments outlined in this chapter were designed to examine the variation in a multi-layered GA. Test 3 - Inducing Phenotypic Variability experiments, use static and dynamic deceptive problems to test the following hypothesis *H₃ Diversity, implicitly maintained by a many-to-one, genotype-phenotype mapping of a GA, implemented by an interpretation of the biological processes of Transcription and Translation, is beneficial in searching noisy and dynamic landscapes.* The experiment results indicated that as problem difficulty increased so too did the benefit of including neutrality and that easier problems seemed harder and harder problems seemed easier, as described by Beaudoin [18]. One reason for the easier problems being more difficult relates to the additional computational overhead associated with neutrality [63]. The results also found that the maintenance of diversity within the population, caused by the dampening of impact from variational operators slowed genetic flow and was beneficial in searching the search space, particularly over changing landscapes.

The results produced exhibit similarities to findings presented in previous work on variation and neutrality [51, 52, 205, 206], all be it using a different representation. The representation presented dictates the effect on genetic flow, which regulates the level of variation and appears to vary according to level of difficulty present in the landscape. In many ways the findings from the experiments are in ways, similar to biological studies carried out by King and Dukes [113] who argue that much of the evolution of proteins is down to neutral mutations and genetic drift. Overall the results of the Test 3 experiments indicate that the hypothesis (H_3) is supported.

The first part of Test 4 - Altering Dimensionality results, examines, over a static unimodal, the hypothesis (H_4) *Introducing neutrality into the representation of a GA, through a modular, fixed non-trivial mapping, is beneficial in searching static landscapes with varying levels of difficulty.* The results of the static unimodal experiments illustrated that as problem difficulty increased, the performance of the neutral representation contained in the MGA proved beneficial in searching the search space and illustrated the MGAs robustness. The findings in relation to variation were similar to past research, in that silent mutation is neutral and doesn't affect the fitness value. However, it does maintain diversity in the population and this can assist exploration [206]. The static unimodal results indicated that stated hypothesis (H_4) is supported.

The second part of Test 4 uses a dynamic unimodal problem to test the hypothesis (H_5) *Including a modular, fixed non-trivial genotype-phenotype map which introduces a more flexible phenotypic structure and a higher degree of phenotypic variation through the use of neutrality, benefits the optimisation of solutions over dynamic landscape problems.* The results can be likened to Wilke et al. [201], who

argued that neutrality provides robustness and that the dynamic unimodal experiments illustrated the robustness of the MGA in the face of a changing landscape. The findings highlighted the benefits associated with the type of neutrality introduced by implementations of transcription and translation, increase as the level of difficulty increases and are more pronounced over a changing landscape. The results concur with previous work showing that the presence of neutral networks offer the MGA the ability to optimise difficult dynamic landscapes [18, 51, 52, 63, 205, 206]. Results from both parts of Test 4 were positive and the stated hypothesis (H_5) was supported.

This chapter examined the levels of variation present in population. The experiments outlined in Test 3 and Test 4 were designed to focus on the impact of the proposed genotype-phenotype map on population variation. The tests were created to develop an insight into why the MGA performed well over more difficult static and dynamic landscapes. For both tests the results indicate that through the level of neutrality contained in the mapping, diversity is implicitly maintained within the MGA population, with the level of diversity maintained at a genotypic level increasing as the level of problem difficulty increased, offering the possibility of increased variability. Having examined the impact of the genotype-phenotype mapping on variation and variability, the next step is to gain a better understanding of the accessibility of phenotypes from genotypes and to further examine the effect of neutrality with regard to the genotype-phenotype mapping. The following chapter, Chapter 7, examines the impact of altering the arity of a multi-layered GA and examines the accessibility of phenotypes from genotypes.

Chapter 7

Examining the Arity of a Multi-layered GA

7.1 Introduction

The aim of this chapter is to examine the arity of the multi-layered genotype-phenotype mapping contained in the MGA, in order to develop an insight into the impact of changing the accessibility of phenotypes from genotypes. To achieve this, a number of experiments were carried out to examine the effect of altering the accessibility and to continue to examine the impact of neutrality in relation to the mapping from genotype to phenotype. The chapter also introduces a variation operator within the layers of the genotype-phenotype map, loosely based on missense mutation as found in nature.

In a SGA, variation is applied solely to genetic structures and not to phenotypic structures and each phenotype is represented by a distinct genotype. However, in

nature we find a layered mapping between the genotype and the phenotype, where a single phenotype can be represented by a number of different genotypes. Also in nature, although most variation happens at the DNA level, mutation can also occur at the RNA level. By introducing a many-to-one genotype-phenotype map and a Missense mutation operator into a GA you can create a basic implementation of these phenomena. Extracts in this chapter are taken from [91]. The chapter is organised as follows: Section 7.2 outlines the process of tuning the MGA's genotype-phenotype map, Section 7.3 gives an overview of the experiments carried out, and Section 7.4 describes the results of Test 5, while Section 7.5 concludes.

7.2 Tuning the MGA Arity

The MGA introduces a tunable, modular, multi-layered genotype-phenotype map, which allows a haploid GA to exhibit, some of the characteristics normally associated with a diploid i.e. a mechanism for allowing alleles or combinations of alleles which proved useful in previous generations [76], this can be viewed as maintaining a form of long term memory, without the need to develop a dominance scheme. The MGA population consists of a population of haploid individuals, which allows for the use of traditional crossover and mutation operators on the genotype. This differs from the approach used by diploid GAs (DGAs) as outlined by Goldberg [76], where each individual has two chromosomes and crossover is divided into two steps and mutation is viewed as being neutral. Another difference between the MGA mapping and that of a DGA, is that in the DGA, a phenotype allele is made up from a single genotype allele which is expressed. In the MGA, a single phenotype allele is made from the

cardinality incorporated in the genotype. In this chapter a number of different cardinalities are examined through the use of different representation mappings. Although the MGA's genotype-phenotype map is non-deterministic, the approach differs from that of real-coded binary representation, which incorporate a gene-strength adjustment mechanism [115]. Real-coded binary representations can use standard crossover operators, but mutation is implicit due to the gene-strength adjustment mechanism [115].

The MGAs used in the dissertation are tuned using three different representation mappings in order to examine the impact of altering the arity: a 4-bit MGA representation; a 6-bit MGA representation and an 8-bit MGA representation. The size of the translation table is determined by the representation chosen. For a 4-bit MGA representation, a translation table of 2^4 is created; with a 6-bit MGA representation, a 2^6 translation table is needed and with an 8-bit MGA representation, a 2^8 translation table is required. The size of the translation table represents the granularity of neutrality which exists within the representation. A 4-bit MGA representation requires 4 bits for each individual element of the phenotype, which we refer to as a *phene*; a 6-bit MGA representation requires 6 bits for each phene and an 8-bit representation requires 8 bits for each phene. The arity controls the granularity contained in the representation mappings, with a 4-bit representation containing a coarser granularity than an 8-bit representation. A form of neighbourhood equivalence [176] is used in the translation phase, to map a predefined number of characters to a given phene, in this dissertation, phenes represent either a 0 or a 1. Figures 4.3, 4.4 and 4.5 in Section 4.4.4 represent a 4-bit MGA representation, a 6-bit MGA representation and

an 8-bit MGA representation, respectively. The Missense mutation operator is outlined in Section 4.4.3, and is operated within the layers of the genotype-phenotype mapping.

7.3 Arity Experiments Overview

The chapter examines the impact of varying the level of neutrality by altering the granularity of the representation, which changes the accessibility of phenotypes from genotypes. In a simplified summary of the biological process, the pathway from DNA to a protein includes a transcription stage, which maps DNA to RNA and a translation stage, which maps RNA to proteins. Altering the arity allows the tuning of the level of neutrality present in the genotype-phenotype mapping. By altering the arity of the genotype-phenotype mapping, the experiments conducted, examine the impact of altering the size of the translation table, as evolution is heavily influenced by neutral mutations and genetic drift [113]. Experiments are also carried out to examine the impact of including a transcription phase in the mapping over more difficult dynamic landscapes. Finally, the experiments attempt to illustrate whether or not, there are benefits associated with the inclusion of variation within the layers of the genotype-phenotype map. Variation within the layers of the genotype-phenotype map is achieved through an implementation of a form of variation operator found in biology, known as *Missense mutation*.¹

The chapter sets out to test the following hypotheses relating to variability in the

¹In biology, a point mutation that changes a codon that normally specifies a particular amino acid into one that codes for a different amino acid is known as a missense mutation.

population:

- (H_7) *Altering the level of granularity, alters the size of the translation table and impacts on the search over more difficult dynamic landscapes.*
- (H_8) *The benefit of including a Transcription phase in a many-to-one, genotype-phenotype mapping, increases as the level of problem difficulty increases over more difficult dynamic landscapes.*
- (H_9) *An interpretation of Missense mutation, included within the layering of a biologically inspired, multi-layered genotype-phenotype mapping GA, assists searching more difficult dynamic landscapes.*

7.4 Test 5 - Altering the Arity

The Test 5 experiments were carried out over a variety of static unimodal OneMax [159] problem landscapes and a number of multi-modal deceptive changing landscapes [73, 87]. The landscapes chosen for experimentation were as follows:

- 30-bit OneMax Problem [159]
- 90-bit OneMax Problem [159]
- 10 3-bit Deceptive Changing Environment [87]
- 30 3-bit Fully Deceptive Deceptive Changing Landscape [87]

In order to examine the accessibility of phenotypes from genotypes, the relationship between altering the granularity of the MGA representation and problem

difficulty are examined, with experiments being carried out over increasingly more difficult landscapes. The OneMax problem experiments include 30-bit and 90-Bit problems, while the deceptive changing landscape experiments use a 10 3-bit loosely ordered fully deceptive changing landscape and a more difficult 30 3-bit loosely ordered fully deceptive changing landscape.

7.4.1 Test 5 Experiment Results

Parameters

The parameters for the experiments are as follows: one-point crossover is used at a rate of 0.7, single-point mutation is used at a rate of $1/l$, where l is the length of the chromosome and missense mutation is at a rate of $5/r$ where r is the length of the RNA string. The population consisted of 200 individuals, with the experiment results averaged over 10 runs. The 30-bit OneMax experiments were run over 50 generations and the 90-Bit OneMax experiments were run over 100 generations. While the 10 3-bit changing deceptive problem and the 30 3-bit changing deceptive problems were run over 5000 and 7000 generations respectively.

30-bit OneMax Problem

Figure 7.1 compares the average-best (off-line) and the average (on-line) performance of an SGA against a number of different MGA representations over a 30-bit OneMax problem. Although there is little difference in the performance of the SGA against that of the variously tuned MGAs, the SGA's performance is marginally better when looking at the off-line results and its on-line performance, is a result of

convergence within the population and the impact of mutation. The MGA results suggest that, similarly to what was reported by Galván-López and Poli [63], that there is little advantage in using neutrality over a relatively easy unimodal OneMax problem.

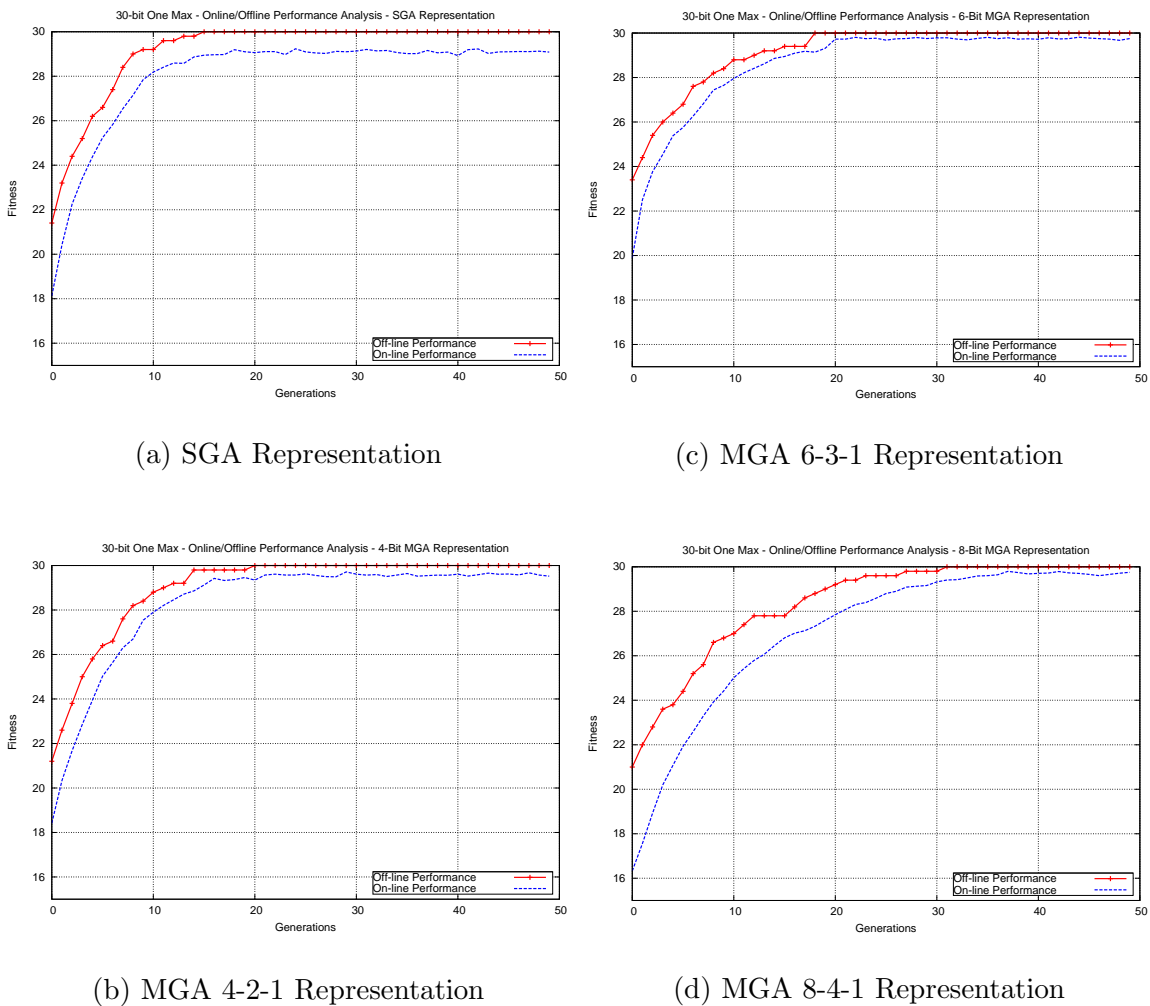


Figure 7.1: 30-bit OneMax Static Landscape

This ties in with work presented by Beaudoin et al. [18] which showed that over less difficult landscapes, the inclusion of neutrality does not assist in searching and

that neutrality makes easier problems harder. Another finding to report from this set of experiments is that the closeness of the online and off-line MGA results, which indicates that there is little exploration taking place (due to the level of problem difficulty associated with the OneMax problem) illustrating that the MGA's representation reduces the impact of variational operators, as reported by Doerr et al. [49]. Also the findings presented in Figure 7.1, differ from Yu and Miller [207], as no positive effect appears to exist from including neutrality over this OneMax landscape.

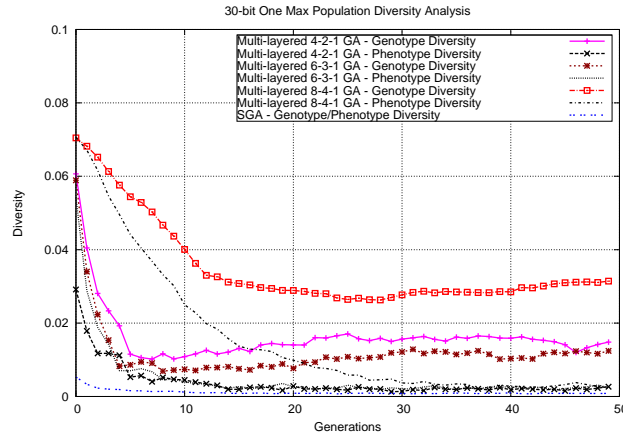


Figure 7.2: Diversity 30-Bit OneMax SGA & MGA

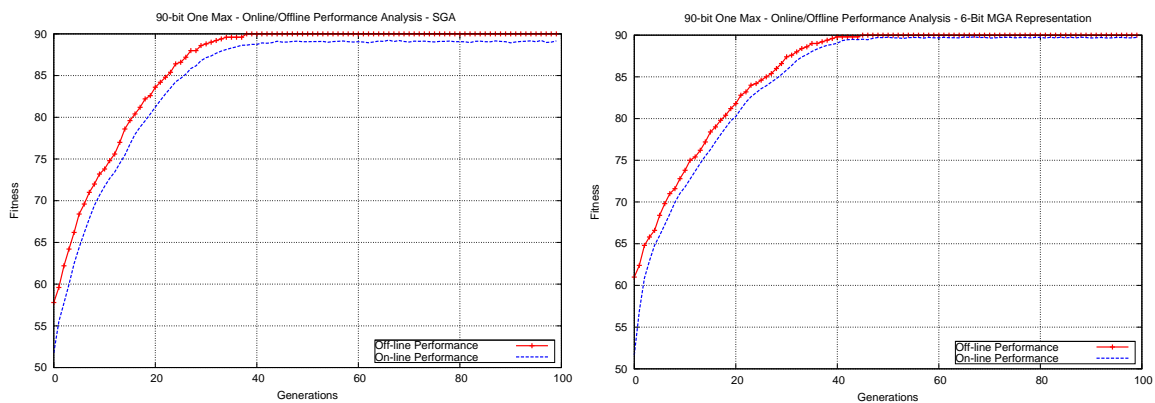
Figure 7.2 contrasts the level of diversity in the population between the SGA and the MGA representations. The SGA phenotypic diversity converges earlier than that of the MGA and there exists a significant difference between the genotypic diversity of the MGA representations to that of the SGA. As the unimodal OneMax landscape was easy to search, the difference in genotypic diversity between both GAs, coupled with closeness of the MGA's off-line and online results would suggest that there are a number of different genotypes sharing the same equivalence neighbourhood and with a redundant representation, a phenotype's phenotypic neighbourhood corresponds to

the phenotypes which are reachable from the given phenotype by a single mutation of the genotype representing it [35]. One conclusion is that the accessibility of phenotypes from genotypes is curtailed, due to the level of difficulty and therefore variability is reduced. To gain a better understanding of the impact of varying the granularity within the representation, the remaining experiments are designed to increase the level of problem difficulty.

90-bit OneMax Problem


By using a 90-bit OneMax static unimodal landscape problem the level of difficulty is increased. Figure 7.3, shows that both the SGA and MGA variations had little difficulty in locating the global optimum. The figures show that given this static unimodal landscape, although level of difficulty has increased from the previous experiment, there is little evidence to argue for the introduction of a neutral representation. However, it appears that over the more difficult OneMax problem, the level of exploration decreases, illustrated by the closeness of the MGA's online and off-line results. In other words, as the level of difficulty increases over the OneMax landscape, the online and off-line results were closer to one another, than those recorded over the easier OneMax landscape. A possible reason for this finding, which was constant as arity increased, is that as reported by Beaudoin et al. [18], harder problems become easier as neutrality is included and the level of exploration reduces accordingly. The on-line performance are marginally better than that of the SGA, as was the case over the previous landscape, illustrating that neutrality provides buffers to dampen the impact of destructive mutations, as identified by Yu and Miller [207].

Again as found by Galván-López and Poli [63] the increase in computational overhead associated with the introduction of neutrality may be out-weighted by the simplicity of the static unimodal landscape. Therefore, for both the unimodal OneMax static landscape used, the level of difficulty doesn't necessitate the introduction of the many-to-one genotype-phenotype landscape associated with the MGA.



(a) SGA Representation

(c) MGA 6-3-1 Representation



(b) MGA 4-2-1 Representation

(d) MGA 8-4-1 Representation

Figure 7.3: 90-bit One Max Static Landscape

Population diversity for the SGA and MGA representations are shown in Figure




Figure 7.4: Diversity 90-Bit OneMax SGA & MGA

7.4. As with diversity in the 30-bit One Max experiments, there is little difference between the SGA and the various MGA representations at the phenotypic level, which is what would be expected, given the low level of problem difficulty associated with the chosen landscapes. However there is a significant difference at the genotypic level, with the MGA variations maintaining a higher level of diversity to that of the SGA, which has converged. This increase in variation combined with the on-line performance, suggests that the many-to-one representation allows genotypic equivalence neighbourhoods emerge.

The results indicate that as the level of problem difficulty increased over the static unimodal OneMax landscape, the number of neutral networks decreased, illustrated by the increase in similarity between the online and off-line MGA results. This ultimately results in a reduction in the accessibility of phenotypes from genotypes. The difficulty with the examination of the impact of arity on variation over unimodal landscapes relates to the level of problem difficulty, as results thus far suggest that there is no advantage in the introduction of neutrality over this class of problem.

In order to further examine the relationship between variation and granularity, the following experiments examine the variability of the MGA over a number of more difficult deceptive changing landscape problems.

10 3-bit Deceptive Changing Environment

Figure 7.5 shows the results of the SGA and the various MGA representations over the 10 3-bit loosely ordered deceptive changing environment. Both the SGA and the various MGA representations, reflecting different levels of granularity, found the global optimum before and after the landscape changed, illustrating that they both managed to adapt and avoided being drawn towards the deceptive attractor. The results indicate that all of the MGA representations out performed the SGA over this landscape, locating the global optimum earlier in the search.

One reason for this is that given the level of difficulty associated with the landscape, the addition of neutrality was beneficial, which is similar to findings by Beaudoin et al.[18]. Although each of the MGA representations chosen out-perform the SGA, what's interesting to note is the differences between the MGA representations with regard to their off-line performance. Results suggest that as arity increases, the online performance improves, indicating an altering of the balance between exploration and exploitation.

The MGA off-line performance dis-improves as the level of arity increases and the on-line performance improves. This finding may be explained by biological studies carried out by Huynen [100] which showed with large amounts of neutrality present in the mapping, neutral paths exist, which allow for smooth exploration. Indicating

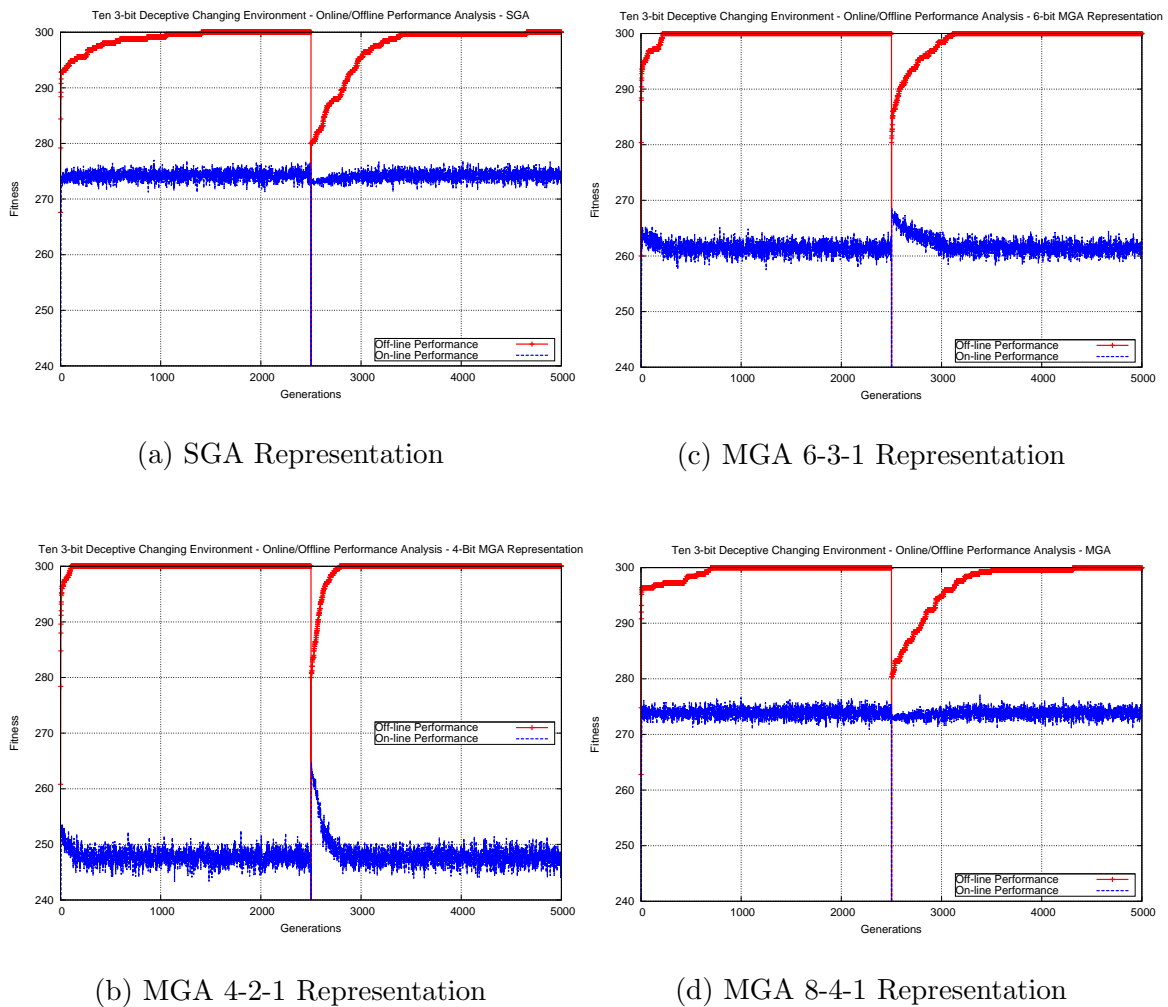


Figure 7.5: 10 3-bit Deceptive Changing Environment

that the balance between exploration and exploitation alters as arity increases. Experiments over a more challenging landscape will examine this further as the balance is critical particularly over changing landscapes. Another finding to emerge from the experiment results is that in comparison with the results reported over the same landscape in Section 5.4.1, there is a significant improvement in optimising the problem. The reason for this relates to the introduction of the missense mutation operator

which increases variability.




Figure 7.6: SGA & MGA Diversity 10 3-Bit Deceptive Changing Environment

Examining the population diversity presented in Figure 7.6, there is a significant difference between the SGA and MGA in relation to variation. Although the phenotypic diversities are relatively similar, there is a greater level of genotypic diversity maintained in the MGA representations, indicating a greater number of different genotypes present in the population. However through the many-to-one representation, many of these genotypes are sharing equivalence neighbourhoods. Due to the nature of the more difficult, multimodal deceptive landscape, in comparison with the unimodal OneMax problems, the MGA on-line performance suggests the level of difficulty leads to an increase in the number of genotypic neighbourhoods being inhabited during the search. Indicating that as the level of problem difficulty increases so too do the number of neutral networks present. The results also appear to suggest that as the arity increases the number of neutral networks decrease, illustrated by the off-line performance and the level of variation.


It appears with this landscape that as arity increases the level of exploration

decreases and the global optimum takes longer to locate. The results also indicate that the level of genotypic diversity maintained by the 8-4-1 MGA representation was less than the MGA variations, due to the increased level of neutrality present in the representation, which as described by Huynen [100], leads to smooth exploration, this may account for its slower performance in locating the global optimum in comparison to the other MGA representations. Another explanation may lie in the increased neutrality associated with an increase in arity, which as reported by Wagner [191] lessens the impact of variation operators, who suggested that neutrality assisted in sheltering a system from the impact of mutation and argued that new adaptations, provided by neutrality, can assist the search. The results outlined in Figures 7.5b, 7.5c and 7.5d indicate that the dampening of the impact of mutation increases as the level of arity increases.


30 3-bit Deceptive Changing Environment

The 30 3-bit loosely ordered fully deceptive changing landscape problem operates on the same principal and increases the level of problem difficulty over the previous 10 3-bit fully deceptive changing landscape. Figure 7.7 compares the performances of the various representations, illustrating that over the more difficult fully deceptive landscape the SGA attempts to locate the global optimum of 900, but fails within the given number of generations, both before and after the landscape changes (see Figure 7.7a). The Figure illustrates the on-line and off-line performance of the SGA and shows that the global optimum is never located before the landscape alters at generation 3000 and once the fitness function has altered, the SGA fails to adapt and


locate the new global optimum.




(a) SGA Representation



(c) MGA 6-3-1 Representation



(b) MGA 4-2-1 Representation



(d) MGA 8-4-1 Representation

Figure 7.7: 30 3-bit Deceptive Changing Environment

In contrast to the SGA results, Figure 7.7 shows an improved performance for each of the MGA representations used. The Figures indicate that the 6-3-1 (Figure 7.7c) and 8-4-1 (Figure 7.7d) representations were successful, both before and after the landscape changed. However, the 4-2-1 MGA representation (Figure 7.7b) failed to recover sufficiently, within the required number of generations, after the changed

environment, indicating that the adaptive quality of the MGA may lessen over more difficult changing landscapes as the level of granularity in the representation decreases.

The results presented in Figure 7.7 show that there is a significant advantage in including neutrality over this class of problem. The results correspond with findings by Ebner et al. [51, 52], who found that high levels of mutation could be sustained by having neutral networks present. They also identified that neutral networks assist in maintaining diversity in the population, which may be advantageous in a changing environment. Also Elgin argued that the random drift caused by neutral variants increases population diversity by expanding the population distribution so as to help find an escape route from local traps [54]. These factors can help in explaining the success of the MGA representations, however the inclusion of a missense mutation operator may also be a factor as it is designed to induce phenotypic variability.

In relation to the off-line performances presented in Figure 7.7, the results would indicate that the SGA representation cannot provide an adequate amount of exploration to succeed over the 30 3-bit loosely ordered deceptive landscape. The MGA representation, in contrast offers the ability, through neutrality, to adjust the balance between exploration and exploitation in order to succeed. However Figure 7.7b shows that the 4-2-1 MGA representation doesn't provide enough exploration to locate the global optimum after the landscape change, within the required number of generations. The other MGA representations (Figure 7.7c and Figure 7.7d) offer an adequate balance between exploration and exploitation to succeed. Again this can be explained through Huynen [100] work, which showed that large amounts of neutrality present in the mapping, provides neutral paths for smooth exploration.




Figure 7.8: SGA & MGA Diversity 30 3-Bit Deceptive Changing Environment

Associated with each of the results conducted over the 30 3-bit fully deceptive changing landscape problem, is the amount of diversity within the various populations. As mentioned previously, as the level of granularity changes, due to changes in the size of the translation tables, so too does the level of variation and the accessibility of phenotypes from genotypes. As with previous results the level of phenotypic diversity maintained by the various GAs remains quite similar (see Figure 7.8). However, the level of genotypic diversity maintained by the MGA representations is again greater than that of the SGA. Over this deceptive changing landscape, the levels of phenotypic diversity maintained by the MGA representations are quite similar to one another. This similarity combined with the online results presented in Figure 7.7, suggests that the number of neutral networks present varies and as arity increases the number of neutral networks increase accordingly, to allow a sufficient amount of exploration.

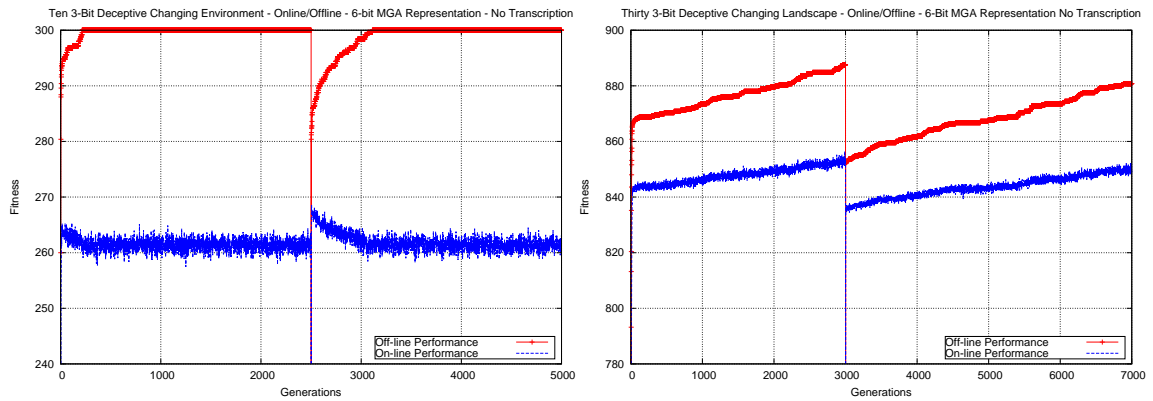
As Nimwegen et al. [143] illustrated, the evolution of the population through neutral networks, suggesting that migrating individuals remain highly connected, that

is with neutral neighbours sharing the same level of fitness. This results in phenotypes which are quite resilient against mutation. When taking neutral theory into account, then the accumulation of neutral mutations should lead to new paths being located [35, 37]. With regard to the SGA, both the genotypic and phenotypic diversity are the same and without an adequate amount of exploration, the population has become trapped on a local optimum.

No Transcription Experiments

To develop an understanding of the impact of the transcription phase on the MGA, experiments were run over the 10 3-bit deceptive changing landscape using a 6-bit MGA representation without transcription, with Figure 7.9 presenting the results. Looking at Figure 7.9a, the results show that although the search locates the global optimum, the performance was similar to the results achieved using transcription. One possible reason is that although there wasn't a changing of characters as would have implemented at the transcription phase, the characters were created directly from the binary string and then passed onto the translation phase without being altered. What the results indicate is that the transcription phase has had very little influence over this particular landscape and that the main forces at work relates to the translation phase and more specifically to the size of the translation table, which sets the level of granularity.

However, over a more difficult landscape, an experiment using a 6-bit MGA representation was run without the transcription phase over the 30 3-bit loosely ordered deceptive changing landscape, illustrated in Figure 7.9b. The results indicate that



(a) 6-3-1 MGA 10 3-bit Deceptive

(b) 6-3-1 MGA 30 3-bit Deceptive Problem

Figure 7.9: No Transcription Changing Environment

the 6-3-1 MGA representation fails to locate the global optimum. This result differs from our previous experiment, shown in Figure 7.9a, suggesting that in the more difficult landscape, the transcription phase assists in the search strategy. One reason for this is that the transcription phase provides a subtle form of variation which may be viewed as being quite similar to inversion, which in turn appears to assist the search strategy over the more difficult deceptive changing landscape.

No Missense Mutation

Figure 7.10 shows the results of not including the Missense mutation operator. Even on the relatively easy 10 3-bit landscape, illustrated in Figure 7.10a, the 6-bit MGA representation fails in its search for the global optimum, indicating that the missense mutation strategy assists in tracking the optimum over a changing landscape and without it the search struggles to escape from a deceptive attractor. These results can also be compared with Section 5.4.1, the difference being the number of

generations over which the problem ran, where missense mutation was included the problem was optimised using fewer function evaluations.

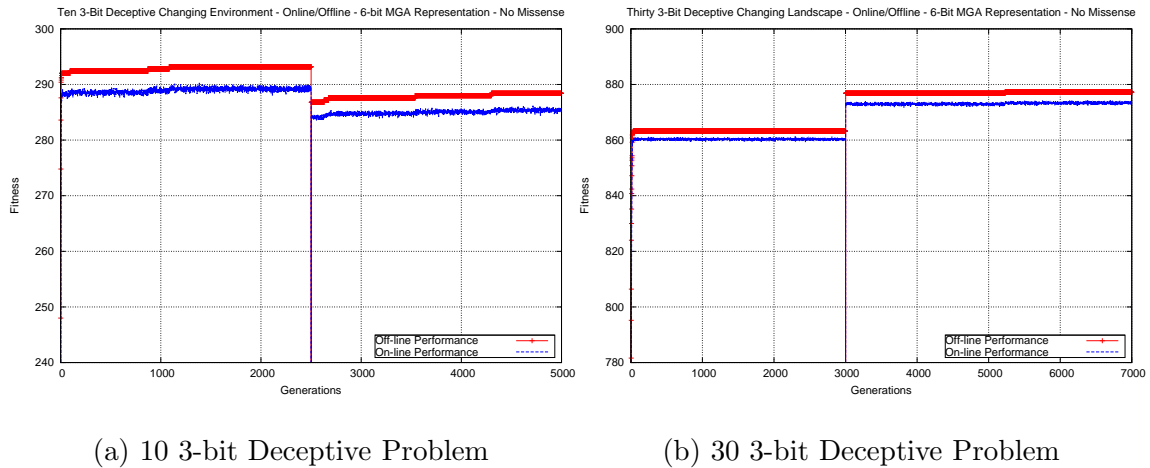



Figure 7.10: 6-3-1 MGA Changing Environment - No Missense Mutation

Figure 7.10b shows that without Missense mutation the 6-3-1 MGA representation fails to escape the deceptive attractor throughout the search. In a similar way to the results shown in Figure 7.10a, there appears to be too little variation and variability, probably due to quite low mutation rates, in the population. Overall, it appears the inclusion of missense mutations in the GA increases the level of variability, thereby allowing the optimisation to take place earlier in the search.

To examine whether the results found by using the multi-layered MGA can be replicated by increasing the level of mutation in the SGA, we ran another set of experiments (see Figure 7.11) where the level of mutation in the SGA was increased to $2/l$, as a method of increasing diversity [32]. The results shown in Figure 7.11a illustrate that in comparison to Figure 7.5a, the extra mutation has improved the performance of the SGA. However, it is worth noting that the off-line performance is



(a) SGA 10 3-bit Deceptive Problem

(b) SGA 30 3-bit Deceptive Problem

Figure 7.11: SGA - Increased Mutation Changing Environment

far lower, a result of the randomness associated with the level of mutation introduced and may prove problematic over a more challenging environment. However, over the more difficult landscape, the SGA fails to locate the global optimum, shown in Figure 7.11b. This appears to indicate that, over the given landscape, even with a higher level of diversity associated with higher mutation, the search strategy of the SGA fails over the given landscape, probably due the randomness associated with high mutation rates.

Statistical Significance

A two-sided paired Wilcoxon signed rank test was carried out on the results of each experiment and were shown to be highly significant ($p < 0.001$) with a p -value $< 2.2e^{-16}$ recorded for each of the experiments illustrated in the chapter.

7.5 Chapter Summary

The aim of this chapter was to examine the arity of the MGA, which alters the size of the translation tables and impacts on the accessibility of phenotypes from genotypes. The chapter also looked at the impact of Transcription and Missense Mutation. The experiments carried out in Test 5 were run over a number of static unimodal landscapes and a number of dynamic multimodal problems. The experiments were designed to test the following hypotheses:

- (H_7) *Altering the level of granularity, alters the size of the translation table and impacts on the search over more difficult dynamic landscapes.*
- (H_8) *The benefit of including a Transcription phase in a many-to-one, genotype-phenotype mapping, increases as the level of problem difficulty increases over more difficult dynamic landscapes.*
- (H_9) *An interpretation of Missense mutation, included within the layering of a biologically inspired, multi-layered genotype-phenotype mapping GA, assists searching more difficult dynamic landscapes.*

Looking at the experiments, designed to test H_7 , conducted over a number of unimodal and multimodal landscapes. The results indicated that there was little benefit in including neutrality on unimodal landscapes and that easier unimodal problems appeared more difficult and more difficult unimodal problems became easier, this trend continued as arity increased. These results are consistent with those found by Beaudoin et al. [18]. The results also illustrated that there little exploration was required due to the level of problem problem and that the MGA's representation

reduces the impact of variational operators, as reported by Doerr et al. [49]. Also as found by Galván-López and Poli [63] the increase in computational overhead associated with the introduction of neutrality out-weights the introduction of neutrality on static unimodal landscape. The results also indicated that as the level of problem difficulty increased over the static unimodal landscape, the number of neutral networks decreased, which ultimately results in a reduction in variability.

Over the more challenging dynamic multimodal problems, the MGA performed better than the SGA and the performance improved as the level of problem difficulty increased. It appears that given the level of difficulty associated with the landscape, the addition of neutrality was beneficial, which is similar to findings by Beaudoin et al. [18]. The results over the multimodal deceptive landscape indicated that as arity increased, the degree of exploration required decreased over the less difficult problem and increased over the more difficult problem, as arity and the number of neutral networks increased. This finding may be explained by biological studies carried out by Huynen [100] which showed with large amounts of neutrality present in the mapping, neutral paths exist, which allow for smooth exploration. Looking at the results of the static unimodal OneMax problem, it appears that as the level of difficulty increases, the off-line performances improve, possibly indicating an increase in the dampening of the effect of variation operators as problem difficulty increases and also as arity increases. This when taken into consideration with the diversity results, also implies that through the many-to-one representation, a number of genotypes are representing the global optimum.

To summarise the H_7 experiments, due to the nature of the more difficult, mul-

timodal deceptive changing landscape, the MGA on-line performance and variation suggests that an increase in the level of difficulty leads to an increase in the number of genotypic neighbourhoods being inhabited during the search, impacting on the evolutionary trajectory. Put another way, with the MGA's neutral representation over more difficult dynamic landscapes, as arity increases, it impacts on the number of neutral networks present, offering paths to escape deceptive attractors, impacting on the search. The findings suggest that over more deceptive dynamic landscapes as difficulty increases, so too do the number of neutral networks and the level of arity impacts on this, thereby impacting on evolutionary trajectory and the search. The reverse is also true, over less difficult dynamic problems as arity increases, altering the granularity, the number of neutral networks decrease, also impacting on the search. Overall the results indicated that the stated hypothesis H_7 is supported.

In relation to the experiments conducted to examine H_8 , the results reported show that in relation to the use of transcription, the advantage of this phase is not apparent over less difficult landscapes as the exclusion of transcription had little or no impact. Once the level of difficulty increased however, as with the more difficult changing deceptive problem, transcription proved vital as without it, the search failed before and after the landscape changed. One possible reason for this is that the transcription phase performs a form of variation, quite similar to inversion and is beneficial as the level of problem difficulty increases. Therefore the hypothesis H_8 can be accepted. Finally, hypothesis H_9 examined the inclusion of an interpretation of missense mutation. Experiments were conducted both with and without missense mutation and results showed that it assisted in the search strategy over both difficult

and less difficult dynamic landscapes, indicating that the use of missense mutation, which operates at levels within the genotype-phenotype map can be beneficial as it induces phenotypic variability. Because of this the hypothesis H_9 was supported.

In order to obtain information on phenotypic variability, an examination of the evolutionary path from the genotypic space to the phenotypic space is conducted. These respective spaces (genotypic and phenotypic) can be viewed as topological spaces which in turn are broken into sets of neighbourhoods. In the genotype topological space, particular genotypes are grouped into neighbourhoods based upon closeness to particular phenotypes. This closeness is not based on similarity but on the accessibility of the phenotype to the genotype topological space and ignores any associated fitness values. The closeness can be determined by the number of mutations away a particular genotype is from a phenotype [59].

Developing an understanding of the impact on the evolutionary path, caused by the multi-layered genotype-phenotype mapping is crucial in investigating the impact of the proposed neutral representation on search. Chapter 8 examines the evolutionary trajectory of both a SGA and a MGA, with both GAs populations visually broken into sets of neighbourhoods based on closeness to phenotypes. The chapter also examines the impact of neutrality on the accessibility of phenotypes from genotypes. The chapter also looks at variation and compares the MGA to a number of SGAs incorporating a number of traditional diversifying techniques. The motivation is to examine how the introduction of Neutral theory into a GA representation affects the evolutionary trajectory of a population and to examine the impact on phenotypic variability and of silent mutations on gene flow.

Chapter 8

Examining the Population

Trajectory & Heterogeneity of a

Multi-layered GA

8.1 Introduction

We can view evolution as operating on the genotype space and when exploring the search space, evolution, through recombination and mutation, defines the search space's variational topology. When discussing variational topology we need to examine both genotypic variational topology and phenotypic variational topology. Nature uses a complex genotype-phenotype map to advance a relatively simple genotype space variational topology to an extremely complex phenotypic variational topology. Toussaint [183] argues that the phenotype space is what should in fact be viewed as the search space for evolution rather than the genotype space. With this in mind,

the genotype-phenotype map is the key to understanding phenotypic variability and it also allows us to gain an insight into how evolution can adapt the search on the phenotype space.

Having established the efficacy of the MGA with Tests 1 and 2, an examination of variation and inducing variability followed using Tests 3 and 4. Results obtained indicated that variation was implicitly maintained, offering the possibility to induce variability which proved beneficial as problem difficulty increased. Test 5 looked at the arity of the MGA and found that altering the granularity of the representation changes the access of phenotypes from genotypes, thereby impacting on the search and suggesting that both problem difficulty and arity, impact on the availability of neutral networks. Test 5 experiments introduced the missense mutation operator, with results illustrating its benefits on the search strategy.

Overall, findings thus far have shown that the proposed representation maintains variation and increases variability. The remaining questions relate to understanding the MGA representations impact on evolutionary trajectory and how it influences variation, accessibility and variability. As population based EAs are often difficult to analyse [48], in order to develop an understanding of how the MGA representation impacts on search, the chapter outlines visually, analyses of genotypic and phenotypic variation and variability by examining the evolutionary trajectory. The chapter then examines heterogeneity, looking at neutral networks and comparing the MGA representation to traditional diversification mechanisms.

This chapter examines the population trajectory and heterogeneity for both the SGA and the MGA. The experiments are separated under two headings, Test 6 and

Test 7. The experiments in Test 6 visually examine the evolutionary trajectory of both the SGA and the MGA through the use of fitness clouds as outlined by Vérel [189]. Test 7 examines Neutral networks, also known as fitness plateaus, which can be viewed as a set of neighbouring solutions with the same fitness value, in other words points in the search space that are connected by neutral point mutations, where the fitness is the same for all points. Test 7 also compares the MGA's ability to maintain heterogeneity and avoid premature convergence with that of the SGA using traditional approaches to dealing with the loss of solutions and sub-solutions.

The chapter contains extracts from [92, 93] and is organised as follows: Section 8.2 gives an overview of the problem landscape chosen for the experiments and describes the fitness plateaus and changing landscape. Section 8.3 gives an overview of the experiments conducted. Section 8.4 examines the evolutionary trajectory of both populations to gain a better understanding of the impact of neutrality on population dynamics over the course of an evolutionary time period. Section 8.5 investigates the impact on heterogeneity caused by the implicit maintenance of diversity due to the layering of the genotype-phenotype mapping of the MGA and Section 8.6 concludes the chapter.

8.2 4-bit Fully Deceptive Landscape

Experiments contained in this chapter were carried out over a 4-bit fully deceptive landscape, as outlined in [198]. The landscape is then altered at a predefined point in the search to produce a changing environment as outlined in Section 2.6.2. This landscape was chosen as it should provide a similar level of difficulty for both the SGA

and the MGA. Figure 8.1 graphically illustrates the 4-bit deceptive problem [198] and a reversed 4-bit deceptive problem. Figure 8.1a, shows the 4-bit problem landscape and Figure 8.1b, the associated heat map, which is designed to illustrate visually the neutral networks (fitness plateaus) contained in the landscape. The landscape change, which occurs at generation 50, is illustrated in Figure 8.1c, which shows the reversed landscape and Figure 8.1d its associated neutral networks and their fitness values.

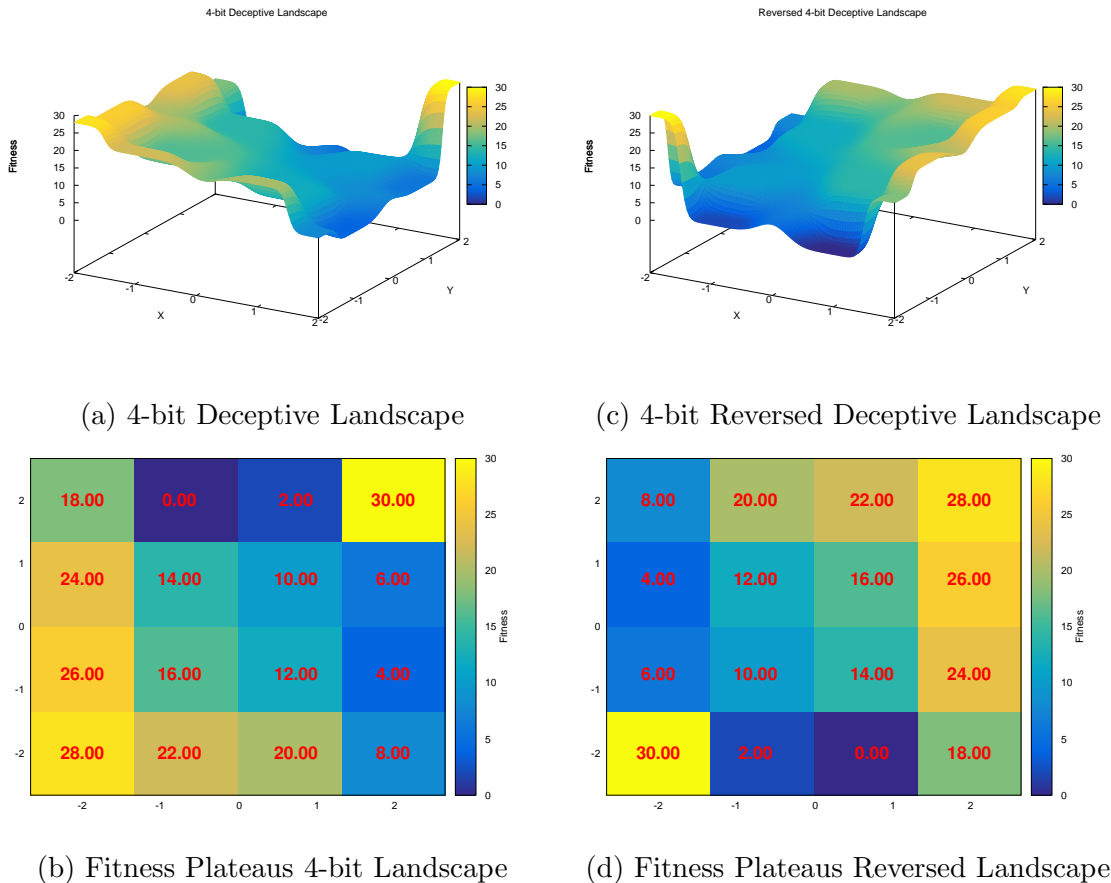


Figure 8.1: 4-bit Deceptive Landscapes & Associated Heat Maps

The results begin by comparing both GAs in a conventional manner based on

performance. Figure 8.2, illustrates the off-line (averaged best fitness) and on-line (averaged fitness) performance for both the SGA and the MGA. The results indicate that the changing 4-bit deceptive landscape initially proved easy for both the SGA and the MGA. However, after the landscape changes at generation 50, the SGA becomes trapped on the local optimum and fails to escape, while the MGA escapes the local optimum and locates the global optimum.




Figure 8.2: SGA & MGA On-line/Off-line Performance Statistics

A Wilcoxon Rank sum test indicates that the results shown in Figure 8.2, that is both the off-line and on-line comparison between the SGA and MGA, were statistically significant. The results indicated that the off-line results for both GAs were highly statistically significant ($p < 0.001$) with a p -value $< 2.2e^{-16}$, similarly the on-line results were also highly statistically significant (p -value $< 2.2e^{-16}$). This difference indicates different dynamics at play in both populations. The remaining analysis of Test 6 examines visually, the evolution of both populations, in an attempt to gain a better understanding of the impact on neutrality on population dynamics over the course of an evolutionary time period and to illustrate the performance of

both population in a dynamic environment. The analysis in Test 7 examines Heterogeneity and compares the MGA to various SGA implementations adopting a number of diversification mechanisms.

8.3 Experiments Overview

8.3.1 Population Trajectory Experiments Overview

The genotype-phenotype map is critical in the operation of phenotypic variability and is essential in allowing an insight into how evolution can adapt the search on the phenotype space. To understand the genotype-phenotype mapping, an examination of population trajectory is required. Test 6 is designed to examine the evolutionary trajectory of both the SGA and the MGA. The populations are presented graphically in order to visualise the impact of layering the representation and introducing neutrality. Test 6 is designed to test the following hypothesis: *H_{10} A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, alters the evolutionary trajectory and is beneficial over a fully deceptive changing landscape.*

8.3.2 Heterogeneity Experiments Overview

Test 7 describes visually the various neutral networks present in the MGA population and illustrates the genotypic equivalence classes, showing that the neighbourhoods and their inhabitants are closely related to the phenotypes they represent, illustrating synonymity, thereby allowing operators to function correctly and ensuring

that the building block hypothesis applies. Test 7 examined the impact of neutrality on population heterogeneity by comparing the MGA to a series of SGA's each incorporating a different diversification mechanism to prevent the premature loss of good solutions and sub-solutions. The approaches used by the SGA included various selection mechanisms; selection noise approaches; selection pressure techniques and operator disruption. Test 7 examines the following hypothesis: *H₆ A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, maintains heterogeneity and is beneficial in promoting exploration over a fully deceptive changing landscape.*

8.4 Test 6 - Evolutionary Trajectory

Parameters

The parameters chosen were as follows: Populations size 20; number of generations 200; number of runs 10; Crossover 0.7; Mutation $1/l$ where l is the length of the chromosome and Missense mutation 0.02. The analysis that follows is broken into three different views of the population trajectory:

- Neutrality & Population Evolution.
- Neutrality & Variation.
- Neutrality & Variability.

8.4.1 Neutrality & Population Evolution

Analysis Before the Landscape Change

In relation to the evolution of the SGA and MGA populations, Figure 8.3 gives an overview of the initial population distribution at generation 0.

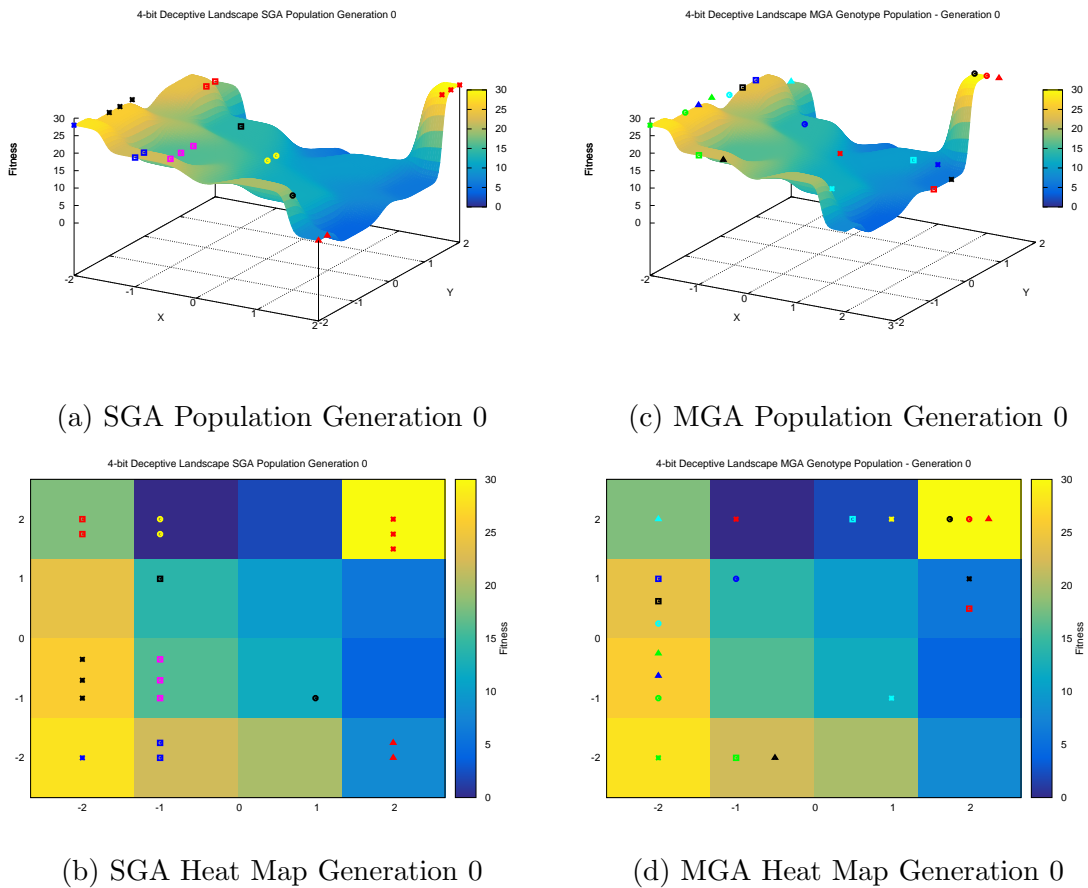


Figure 8.3: SGA & MGA Generation 0

Figure 8.3a and the associated neutral networks Figure 8.3b, illustrate that the initial SGA population of 20 individuals is randomly spread over the landscape. Figure 8.3c shows the MGA population distribution over the problem landscape at generation 0 and its neutral networks in Figure 8.3d. Each of the coloured shapes represent

an individual in the population, with identical genotypes, represented by the same colour and shape, residing in an equivalence neighbourhood. Different genotypes representing the same phenotype reside on the same plateau, but are illustrated using different shapes or colours. Figure 8.3a and Figure 8.3c illustrate that the initial population distributions for both algorithms are quite similar and is emphasised in Figure 8.3b and Figure 8.3d, with different individuals represented by different colours and shapes, identifying different genotypes sharing similar fitness plateaus or neighbourhoods. One of the main differences between both GAs is that with the SGA, there is a one to one (1 : 1) representation between the genotype and the phenotype. The MGA incorporates a many-to-one ($M : 1$) representation, which manifests itself in a number of different genotypes representing the same phenotype.

Examining the populations for both GAs at generation 50 (Figure 8.4), which is the last generation before the landscape changes, the SGA's population has converged, apart from the impact of mutation (Figure 8.4a and Figure 8.4b). A large portion of the MGA's population has also located the global optimum, however, due to the genotype-phenotype mapping, genetic drift slows and the population doesn't converge [10, 54]. Figure 8.4c and Figure 8.4d, show the many-to-one representation present in the MGA population.

The MGA population's evolutionary trajectory differs considerably by increasing variation and avoiding converging on the global optimum. Therefore the population consists of a number of individuals located on different neutral networks, which is a result of gene flow due to the presence of neutrality. Work by Cohoon et al. [33], drew on the idea of *punctuated equilibria* and concluded that by having equivalent

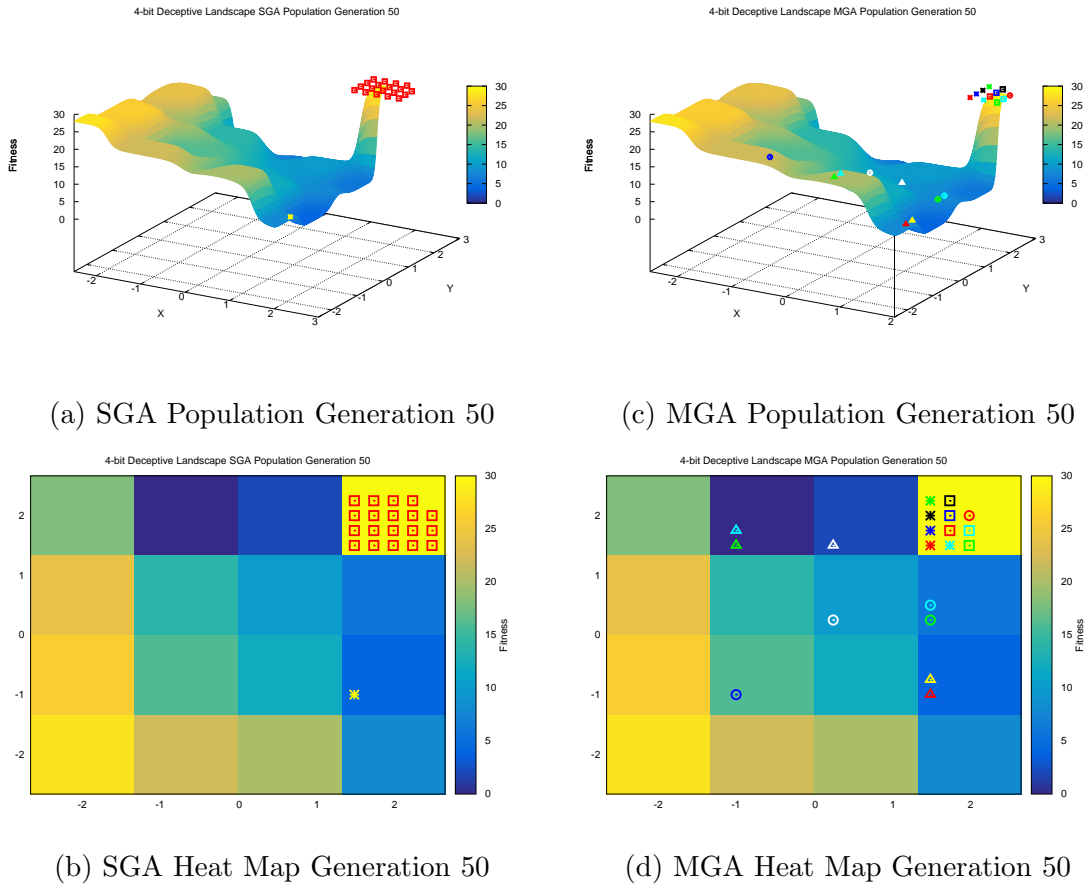


Figure 8.4: SGA & MGA Generation 50

classes a many-to-one representation exists, which proved beneficial to the search. It's interesting to note the location of the populated neutral networks in relation to variability, which as a result of connectivity associated with the MGA representation are, in the main, located in neighbouring plateaus.

Analysis After the Landscape Change

Figure 8.5, illustrates the SGA and MGA population distribution when the landscape changes at generation 51. Figure 8.5a and Figure 8.5b shows the SGA pop-

ulation, which now converges on the local optimum as the landscape has reversed in relation to the fitness function. It also indicates that all but one member of the population, as a result of convergence, are located on the local optimum (formally the global optima).

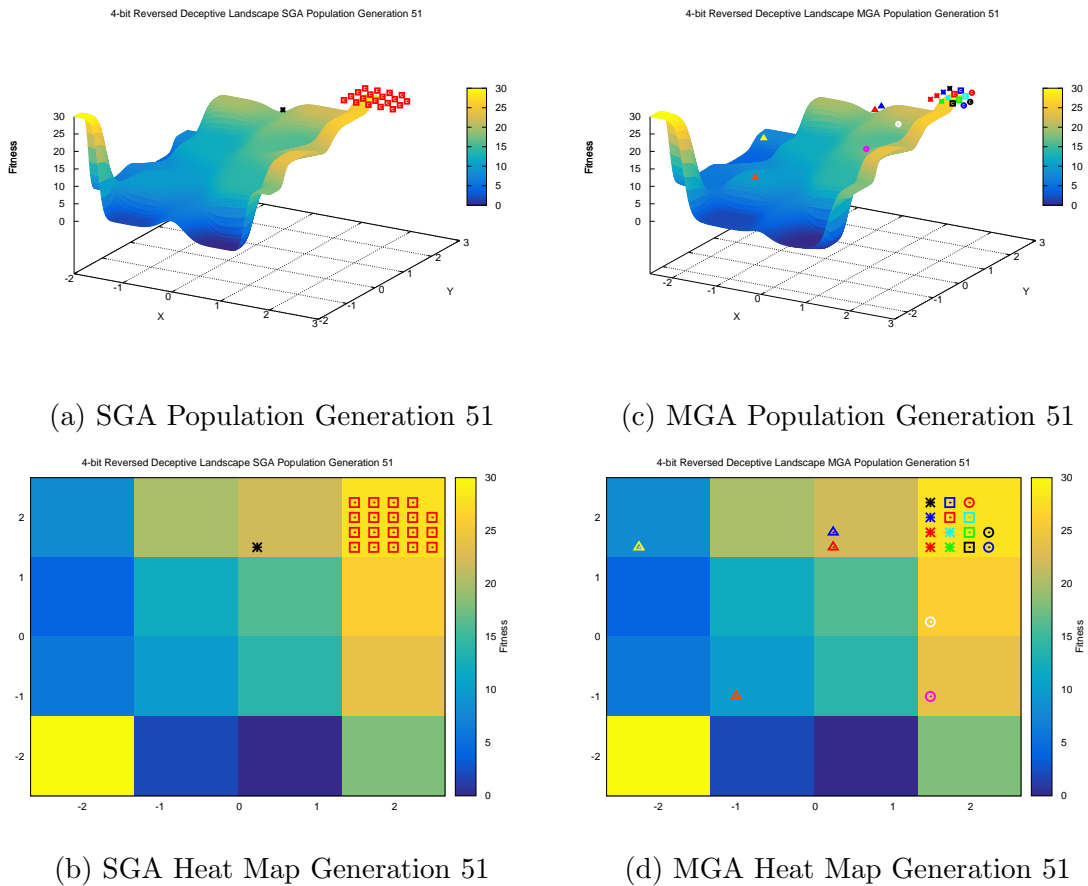


Figure 8.5: SGA & MGA Generation 51

The MGA population has a wider distribution on the problem landscape (see Figure 8.5c), with the bulk of the population located on the local optimum. Figure 8.5d indicates that due to the presence of neutrality in the representation, the population is dispersed over a wider number of fitness plateaus. As found by Elgin [54]

random drift caused by neutral variants increases population diversity by expanding the population distribution, this in turn assists in finding an escape route from local traps.

Looking at the population evolution through to generation 200 for both GAs, as seen in Figure 8.6, the SGA population remains trapped on the local optimum (see Figure 8.6a and Figure 8.6b), while the MGA population has escaped and is now clustered around the global optimum and neighbouring plateaus (see Figure 8.6c and Figure 8.6d).

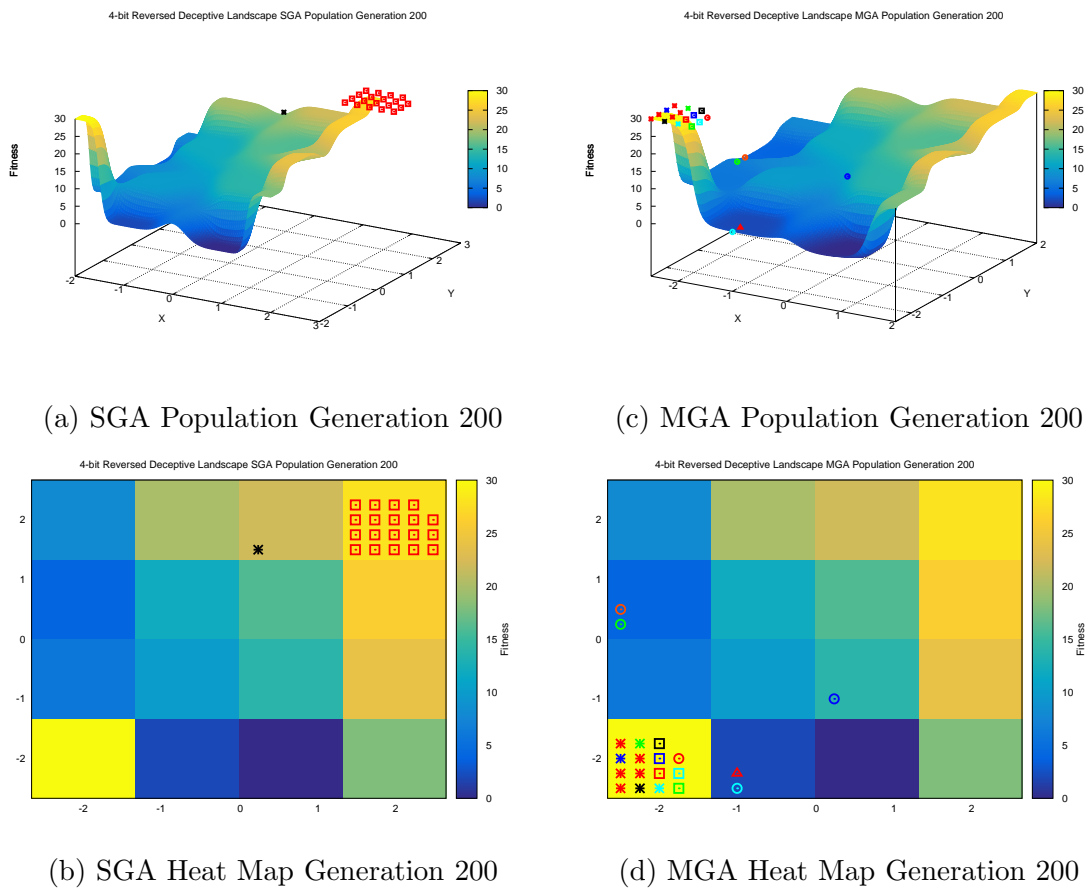


Figure 8.6: SGA & MGA Generation 200

This appears to indicate that the MGA, through the genotype-phenotype mapping, implicitly maintains a level of genetic diversity within the population and is resistant to convergence on a single neutral network as a result of increased variation and variability, thereby offering the ability to induce variability and adapt in a changing environment. These results agree with Shipman [167], who found neutrality to be advantageous where neutral networks are distributed over the search space with a high degree of connectivity between them. Also, as the representation was designed to be uniform and synonymous [154], the results found illustrate that the locality of reference and the connectivity are both high. These findings suggest a similarity to work by Correia [35], who argues that there are synonymously redundant representations that allow connectivity to be increased between phenotypes, when compared to non-redundant representations. The findings are also in line with Nimwegen et al. [143] who illustrated the evolution of the population through neutral networks, suggesting that migrating individuals remain highly connected, that is with neutral neighbours sharing the same level of fitness. This results in phenotypes which are quite resilient against mutation. The next set of figures, illustrate graphically the variation topology.

8.4.2 Neutrality & Variation

The Figures presented under Neutrality and Variation are designed to graph the Hamming difference between individuals in the population in relation to all other population members in order to visually examine the variation for each generation. As a traditional SGA maps directly from the genotype to the phenotype, both the

genotypic search space and the phenotypic search space are identical. With the MGA mapping, neutrality increases the genotypic search space in comparison to the phenotypic search space.

Using normalised Hamming distances between individuals, Figure 8.7 illustrates the population variation. Figure 8.7a shows the genotype/phenotype diversity for the SGA, while Figure 8.7b, illustrates the genotypic diversity for the MGA and Figure 8.7c, the phenotypic diversity at generation 0. Figure 8.7 shows the overall level of diversity, with the generation under observation highlighted in yellow (Figure 8.7d).

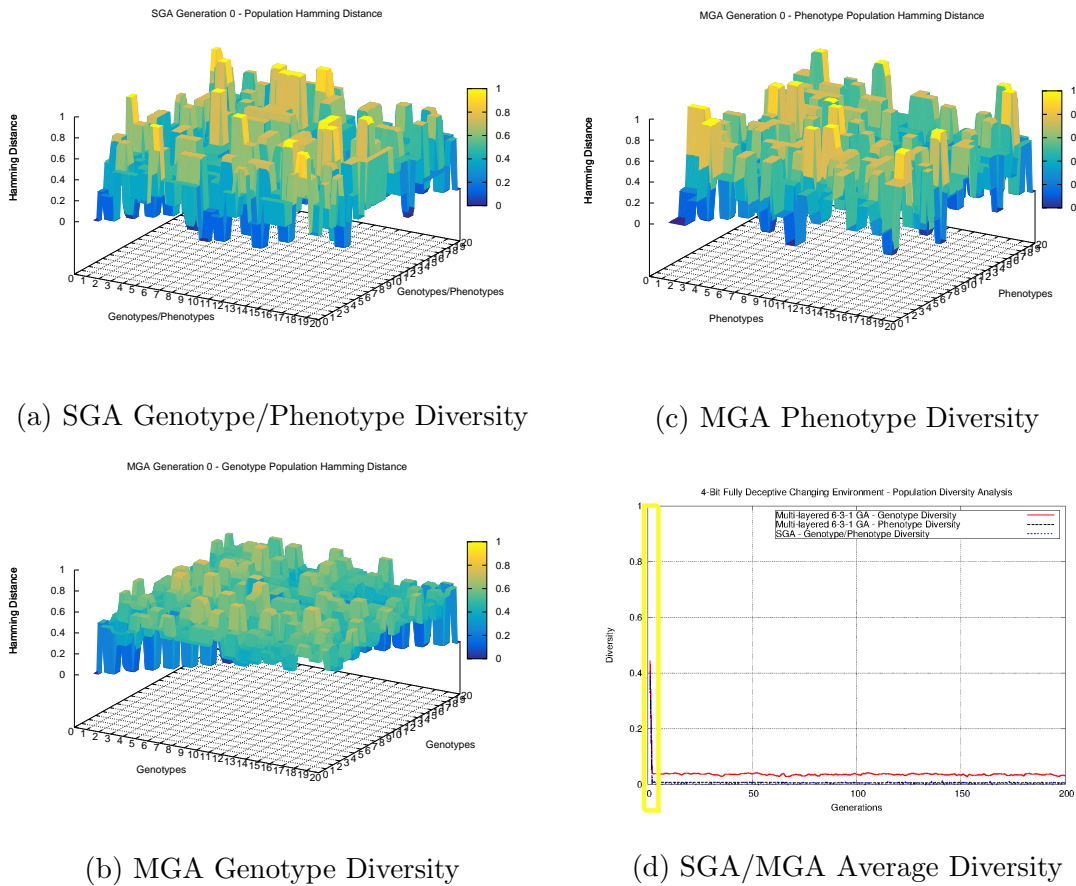


Figure 8.7: SGA & MGA Diversity - Generation 0

From the figures, it appears that at generation 0, there is quite a large level of diversity present in both populations. This can be explained as evolution has not yet begun and the populations have been randomly generated. Also, as neutrality is introduced the MGA's genotypic search space increases, while the MGA's phenotypic space (Figure 8.7c) is quite similar to that of the SGA.

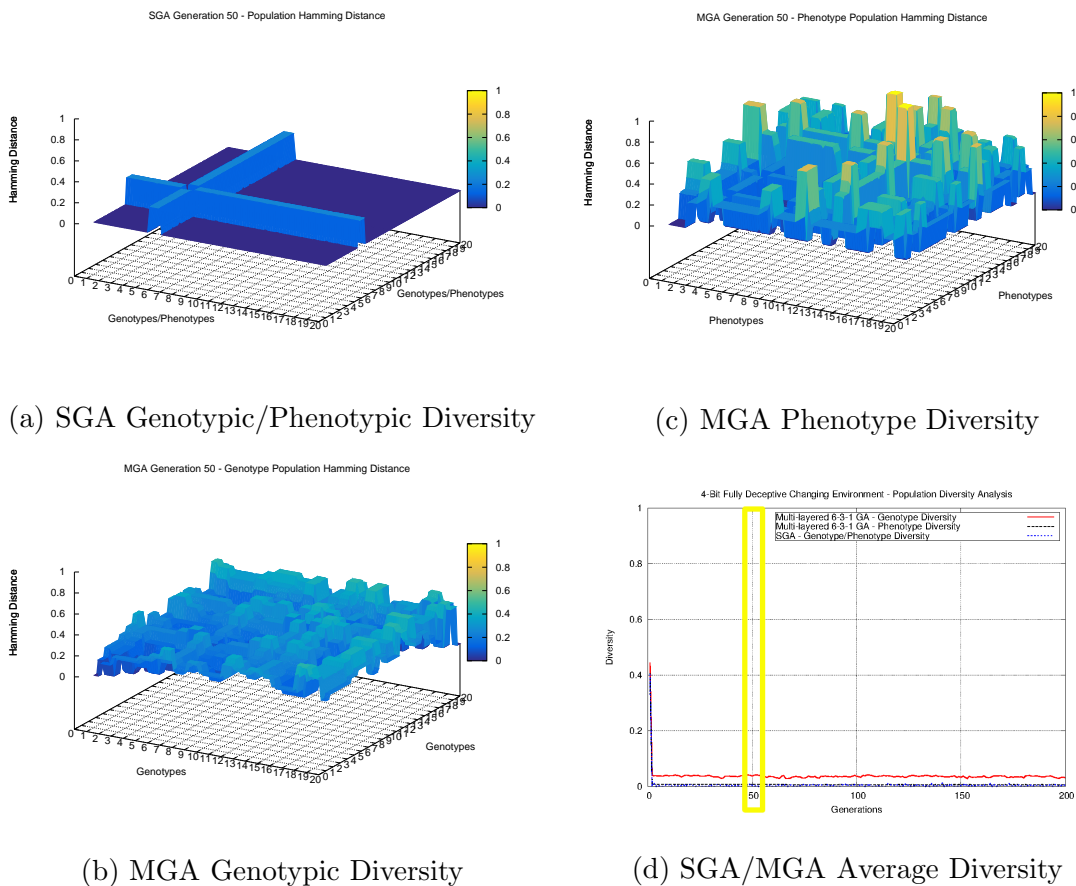


Figure 8.8: SGA & MGA Diversity - Generation 50

As the populations evolve to generation 50, convergence has occurred in the SGA (apart from the influence of mutation) and the population is shown in Figure 8.8a, with the relatively smooth plot, illustrating the lack of variation present in the popu-

lation. In comparison the MGA population preserves variation and genetic diversity is maintained within the genotypic search space. This is shown in Figure 8.8b, which illustrates a more rugged plot compared to that of the SGA, indicating a higher level of variation associated with the neutral mapping. The MGA phenotypic search space (Figure 8.8c) presents the diversity at a phenotypic level, again the plot shows a more rugged landscape illustrating the phenotypic variation and the avoidance of convergence. Figure 8.8d shows the overall level of variation during a run, with the generation under observation being highlighted in yellow.

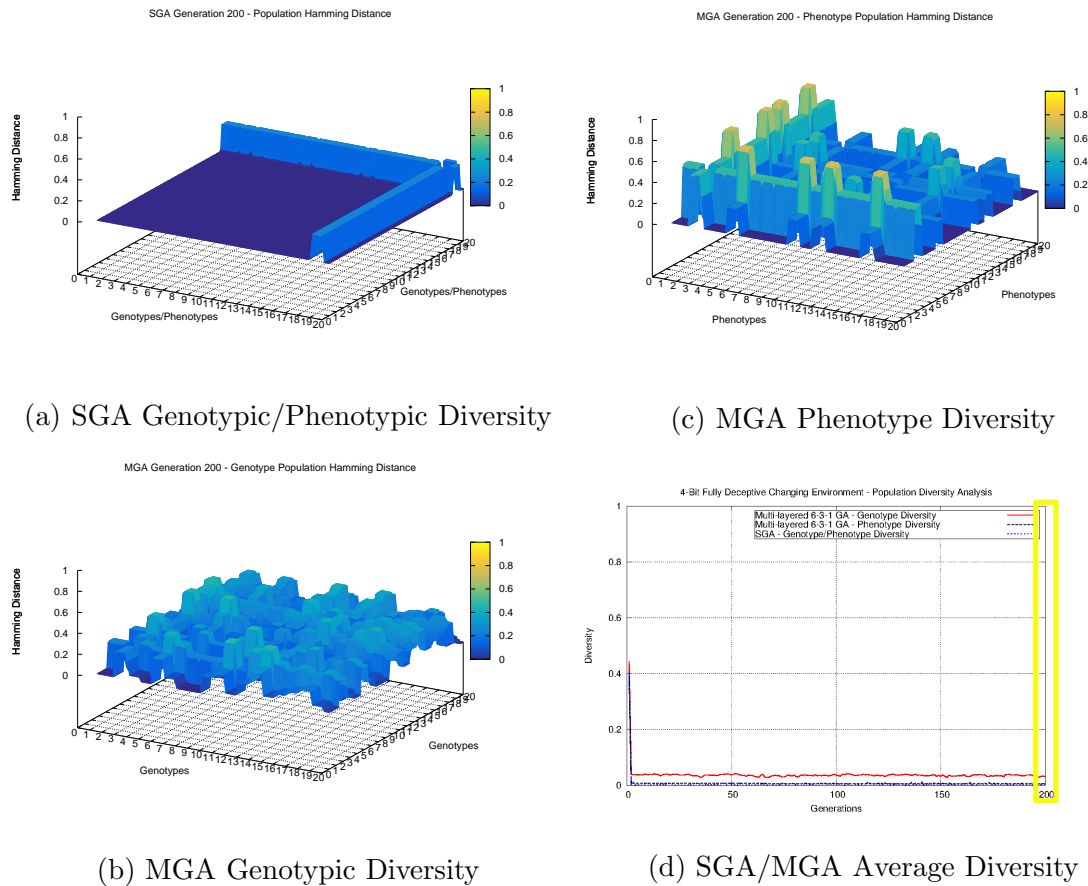


Figure 8.9: SGA & MGA Diversity - Generation 200

In Figure 8.9 we see the variation at generation 200 (Figure 8.9d), the SGA population remains trapped on the local optimum and contains little diversity to find paths to allow it to escape. This situation is shown in Figure 8.9a, reflecting the lack of variation within the population when viewed as a function of Hamming distance. The MGA genotypic variation and phenotypic variation at generation 200 are shown in Figures 8.9b and 8.9c respectively. The level of phenotypic diversity present, relates to the many-to-one genotype-phenotype mapping and to individuals of the population being part of other neutral networks, meaning the population has not converged. The results on neutrality and variation presented illustrate graphically the impact of the MGA representation on genotype and phenotype variation during an evolutionary run. The results are similar to those reported by Yu and Miller [205, 206] and Ebner et al. [51, 52], in relation to neutrality and the maintenance of diversity. Having examined the impact of neutrality on variation we now examine the impact of neutrality on phenotypic variability.

8.4.3 Neutrality & Variability

The final examination of the effects of neutrality on the evolutionary trajectory of the MGA population relates to the impact of the neutrality on variability. In relation to describing a genotypic neighbourhood, by using single point mutation with a one-to-one genotype-phenotype map, then the number of reachable genotypes from a given genotype is $(A - 1)L$, where A is the number of alleles available and L is the length of the genotype [109]. In order to examine the effect of single bit mutation, this process is repeated for both the SGA and MGA with an individual being randomly selected

from both the SGA and MGA populations from each generation. All of the bits are then flipped one bit at a time in sequence and the impact on fitness is noted. The aim of this approach is to examine how the presence of neutrality impacts on the mutation operator and variability, as a genotype in the MGA population, when mutated can produce either a silent or an adaptive mutation.

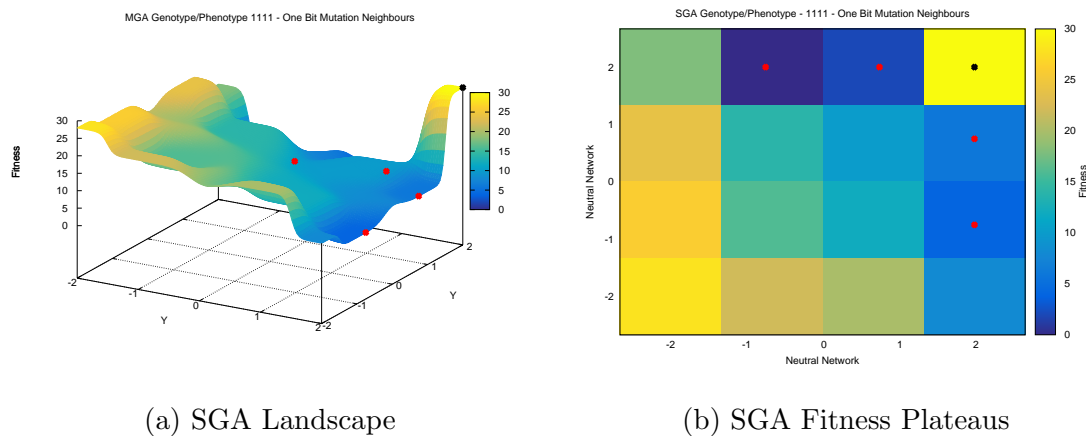


Figure 8.10: Single-bit Mutation SGA Genotype/Phenotype 1111

Figure 8.10, illustrates the impact of mutation on each allele of the genotype 1111. The initial genotype is shown in black with each of the genotypes created by a single bit mutation shown in red, depicted in Figure 8.10a. As the SGA has a one-to-one genotype-phenotype mapping the result of each single bit mutation causes the creation of a different genotype which represents a different phenotype, located on a different fitness plateau, see Figure 8.10b.

Mutation impacts the MGA population in a different as there exists a many-to-one mapping between the genotype and the phenotype which increases variation and variability. Looking at Figure 8.11, which represents the outcome of a single bit mutation for each bit of the genotype represented in black and the mutated genotypes in red

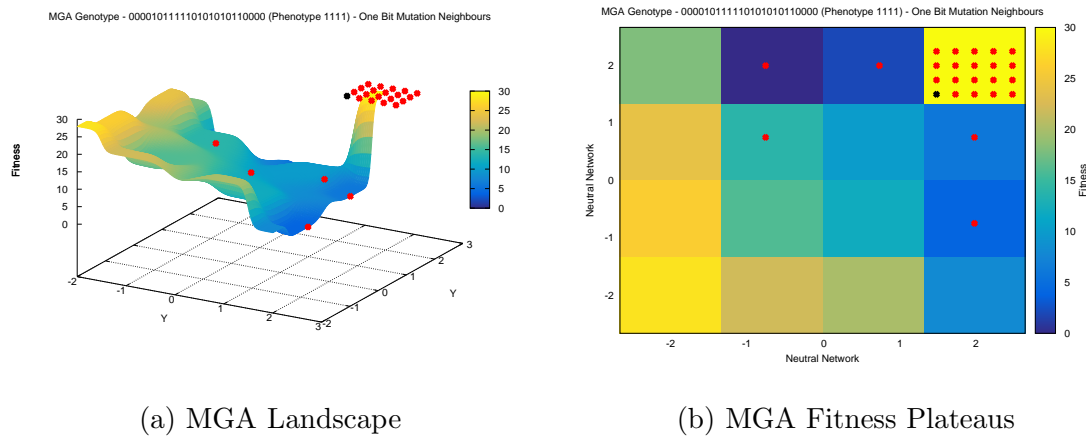


Figure 8.11: Mutation MGA Genotype 00001011111010101010000 Phenotype 1111

(Figure 8.11a), we can see that of the 24 mutations, one for each bit in the genotype, there were 5 adaptive mutations, illustrating the increase in connectivity. By changing the primary structure, the MGA possesses the ability to increase both variation and variability. The silent mutations are located in the same neutral network and are shown in Figure 8.11b. This is as a direct result of the presence of neutrality in the genotype-phenotype map, due to the inclusion of an implementation of transcription, translation and dynamic activation. The findings show the influence of a non-trivial neutral representation, where different genotypes in a neutral set allow for different phenotypic distributions [176, 183].

Missense Mutation

The MGA representation includes missense mutation which operates within the layers of the genotype-phenotype mapping. The missense operator mutates single characters and can be silent or adaptive (see Figure 8.12). Figure 8.12a and Figure 8.12b illustrates the phenotypic distribution for individual when an adaptive

missense mutation takes place following an adaptive single bit mutation, illustrating that the combination of single-bit and missense mutation increases variability. This section highlighted the effect of both mutation and missense mutation, with results illustrating that the SGA can only access neighbouring local plateaus, while the MGA, where mutation can be phenotypically silent or adaptive, provides greater variability [183].

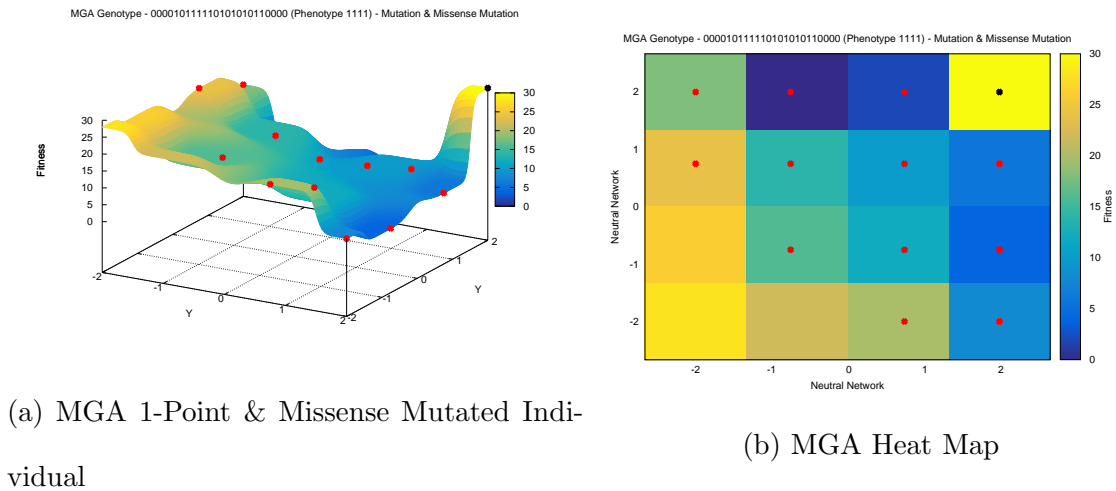


Figure 8.12: MGA 1-Point & Missense Mutated Individual

The Test 6 findings show visually an interpretation of Toussaint’s [184] idea of non-trivial neutrality, where different genotypes in a neutral set allow for different phenotypic distributions and how evolution can adapt the search on the phenotype space [183]. Although two genotypes can be considered equivalent if they represent the same phenotype, however their variational topologies can differ [176]. But ultimately the genotype-phenotype map “*induces a variational topology on the phenotype space depending on the topology of the genotype space*” [183]. The results of Test 6 illustrate graphically the variation and variability associated with the MGA. The find-

ings here suggest that through the neutral representation and dynamic activation, the representation which is, designed to be uniform and synonymous, allows connectivity to increase, as found by Correia [35] and shines some light on the interrelationship between neutrality and connectivity. A further question relates to the heterogeneity and as this dissertation suggests that the representation presented maintains variation, thereby preventing the loss of good solutions, and increases variability through increased connectivity, then how does this compare with traditional diversification methods designed to prevent the loss of valuable genetic material. This is the subject of Test 7 which examines heterogeneity.

8.5 Test 7 - Heterogeneity

Test 7 examines the heterogeneity resulting from the implicit maintenance of diversity within a population through the inclusion of a layered genotype-phenotype map based on the principals of Neutral theory [111], comparing it with traditional techniques adopted by the SGA to prevent the loss of good solutions or sub-solutions. Test 7 is designed to test the following hypothesis H_6 : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, maintains heterogeneity and is beneficial in promoting exploration over a fully deceptive changing landscape.*

In simple GAs (SGA) the selection mechanisms allow the propagation of fitter solutions at the expense of less fit solutions. This in turn leads to the convergence of the population. Convergence will also occur even where you have alternative solutions with identical fitness. Solutions and sub-solutions are normally lost in a SGA pop-

ulation for three reasons: *selection pressure*, *selection noise* and *operator disruption*. Selection pressure occurs as the result of the selection process, with less fit solutions disappearing from the population. Selection noise is the result of the variance of the selection process due to random choices between identically fit solutions. Operator disruption takes place through the implementation of the crossover and mutation operators which possess the ability to destroy good solutions. In an attempt to avoid premature convergence, techniques for diversifying a population generally attempt to reduce individually or in combination with one another, selection pressure, selection noise or operator disruption [126].

Variation within a population can serve a number of purposes, such as delaying convergence in order to promote exploration, and the location of multiple solutions. Because of this the maintenance of diversity within a population is a desirable feature for GAs. Diversification methods capable of reducing all three criteria, selection noise, selection pressure and operator disruption exist. The problem lies in the fact that reducing all three criteria to arbitrarily low levels results in the GA carrying out little or no useful search [126]. The maintenance of diversity for its own sake is undesirable, what is required is diversity that promotes good strings [74]. Past research into diversity and methods to maintain diversity within the population include De Jong [43, 135], Whitley [196], Eshelman [55], Goldberg [71], Grefenstette [79, 32], Mahfoud [126, 125], Bickle [21], Morrison [134], McPhee [129], BarkerbarkerDiversity00, Hutter [99], Motoki [137] and Doerr [48] to name but a few.

By comparing the performance of the MGA with a SGA containing a number of diversifying techniques, we can evaluate the diversity maintenance mechanism implicit

in the MGA representation. With regard to the three mechanisms for promoting diversity, theory would suggest that Stochastic Universal Selection (SUS) offers the ability to minimise selection noise, fitness scaling can decrease selection pressure and lower rates of crossover and mutation can reduce operator disruption [126]. However it is worth remembering that a relationship exists between convergence and diversity in a GA population. If there is no operator disruption and no selection pressure, then the GA will maintain its initial population and won't perform any meaningful search. Also, with regard the mutation operator, as mutation rates increase, the diversity produces is usually not useful as the GA is approaching random search.

8.5.1 Test 7 Experiments

Parameters

The experiments were conducted over the 4-bit fully deceptive changing landscape as outlined above and the parameters chosen were as follows: Populations size 20; number of generations 200; number of runs 10; Crossover 0.7; Mutation $1/l$ where l is the length of the chromosome; Missense mutation 0.02. For the heterogeneity experiments various selection mechanisms were used in conjunction with the SGA and tournament selection was used for the MGA, with a tournament size of 4. Various scaling mechanisms were used with the SGA and no scaling mechanisms were included in the MGA. Over the set of experiments designed to examine various diversifying techniques, the selection mechanisms and scaling methods of the SGA were altered. A number of experiments were also conducted where the SGA's variation operators were increased. The motivation for these changes is to vary the selection noise, selection

pressure and operator disruption for the SGA, thereby examining the impact of the inclusion of well understood diversifying techniques into the SGA, which in turn are then compared with the performance of the MGA.

The experiments conducted for Test 7 can be split into three categories:

- The *selection noise experiments* explore the use of Stochastic Remainder Selection and Stochastic Universal Selection.
- The *selection pressure experiments* examine, Linear, Window, Sigma Truncation and Boltzmann scaling techniques and niching techniques, such as, Crowding and Incest Reduction.
- The *operator disruption experiments* look at the impact of increasing rates of crossover and mutation.

8.5.2 Neutral Networks & Neighbourhood Equivalence

The neutral networks representing various fitness values are shown in Figure 8.13. Neutral networks can be viewed as genotypic neighbourhoods where genotypes representing the same phenotypes group together on the same fitness plateaus. Before the landscape changes at generation 50, the most prominent neutral network represents the global optimum phenotype (1111). As evolution progresses, the MGA, through its $M : 1$ representation, allows the size of neutral networks to adapt as the population evolves. The next largest neutral networks represent the phenotypes (0111), (1110), (1101) and (1011), which are the four genotypes closest in Hamming distance to the optimum, indicating that the MGA's population evolves towards neighbouring

neutral sets.




Figure 8.13: Neutral Networks

As the landscape changes the largest neutral network represents the new local optimum (1111), as this was the global optimum at the time of the change, As the search continues the MGA’s population escapes the local optimum and has located the global optimum. At this point, the most prominent neutral network represents the phenotype (0000) (the new global optimum). Furthermore, the MGA population has evolved towards a different group of neutral sets representing the phenotypes (0001), (0010), (0100) and (1000), all close, phenotypically, in Hamming distance to the global optimum. Examining the evolutionary trajectory from the phenotype space illustrates that the MGA’s representation maintains good solutions and sub-solutions within the population allowing the processing of good schemata. As the Building Block hypothesis argues, the power of a GA lies in being able to find good *building blocks* [67]. Viewing the search space as argued by Toussaint [183] allows an insight to be developed into how evolution can adapt to the search.

Figure 8.14 looks at the make-up of the neutral sets, examining the number of




Figure 8.14: Neutral Networks - Identical Genotype

identical genotypes in each set. The figure indicates that there is a high degree of variation maintained within each neutral set, with low numbers of identical genotypes present. The number of identical genotypes representing the global optimum varying between 2 and 4, illustrating the impact of the $M : 1$ representation on the population. Overall, the results indicate that the representation, implicitly maintains useful building blocks within the population, which promotes good strings and assist in improving the adaptability of the MGA. What this translates to is a situation where the variation in the population contains genotypes which are close to one another in terms of the phenotypes they represent rather than their genetic makeup. To summarise, the Figures show that the neutral networks are constructed by combining identical phenotypic structures as reported by Reidys et al. [150] and illustrate that the most prominent neutral networks have a high degree of connectivity and as argued by Shipman [167], neutrality is advantageous where neutral networks are distributed over the search space with a high degree of connectivity between them. These results also align with research into random walks on graphs [124, 144] illustrating that

nodes within a network with the highest connectivity tend to be visited most often [11]. Thereby allowing the prediction that search within a neutral network will not be a pure random drift, instead it generally has a bias and concentrates on individuals of the network where connectivity is highest [11]. The results shown appear consistent with this as high levels of connectivity as rewarded with frequent visits, as shown in Figures 8.13 and 8.14.

8.5.3 Selection Mechanisms

Roulette Wheel Selection

Roulette Wheel Selection offers each individual a roulette wheel slot in proportion to the fitness of the individual. This results in the creation of a biased roulette wheel which proportionally favours the fittest individuals in the population. Figure 8.15 illustrates both the on-line and off-line performances for the SGA using roulette wheel selection and MGA. The SGA locates the global optimum early in the search but fails after the environment changes and remains trapped on the local optimum, shown by the off-line performance. Roulette wheel selection is highly noisy and assists in the removal of diversity from the population. Comparing the on-line results of both algorithms we can see that the MGA maintains a greater degree of diversity within the population. The diversity maintained by the MGA contains alleles of equal fitness which assists in diversity preservation as alleles of different fitness selection pressure lead to fast convergence [71] and small population selection schemes are subject to the vagaries of genetic drift [75].




Figure 8.15: SGA (Roulette Wheel) & MGA On-line/Off-line

Tournament Selection

A Tournament selection mechanism randomly chooses a predefined number of individuals from the population, who in turn compete amongst one another and the fittest individual goes forward for crossover. Selection pressure can be adjusted by altering the tournament size. A large tournament reduces the probability of weak individuals being selected.




Figure 8.16: SGA (Tournament) & MGA On-line/Off-line

Figure 8.16 shows the success of the MGA in locating the global optimum and the maintenance of useful diversity within the population. The SGA locates the global optimum before the landscape change but cannot escape from the local optimum after the environment changes.

Ranking Selection

With Ranking Selection the population are ranked according to fitness value. Individuals are then assigned an offspring count as a function of their rank. A drawback of this method is the uncoupling of the fitness function from the underlying objective function [67]. Figure 8.17 compares the performance of an SGA incorporating a ranking selection mechanism.




Figure 8.17: SGA (Ranking) & MGA On-line/Off-line

The results indicate an element of randomness in the use of ranking selection, illustrated by the SGA on-line performance, which shows a large variance within the population. Again the MGA locates the global optimum before and after the landscape change. Although the SGA locates the global optimum both before and

after the environment change it also loses it, due mainly to the degree of randomness present as a result of ranking selection.

Figure 8.18 outlines the degree of convergence within the MGA population and the SGA population using Roulette Wheel, Ranking and Tournament selection mechanisms. The results indicate that using Tournament selection resulted in the SGA population converging very quickly (Generation 5). Roulette Wheel selection had similar results (Generation 9). Ranking selection produced very different results and maintained a degree of variation within the population. However, the variation maintained was based on randomness and was not useful, illustrated by the search continuously moving off the global optimum. The MGA (using Tournament selection), produced the best performance by maintaining useful variation both at genotypic and phenotypic level.




Figure 8.18: SGA & MGA Convergence Rate

8.5.4 Selection Noise

As outlined by De Jong [43], the variance of selection is one of the main contributors to the idea of convergence. In order to examine selection variance or selection noise, we compare a number of selection mechanisms designed to reduce selection noise, namely Stochastic Remainder Selection (SRS) and Stochastic Universal Selection (SUS). With Stochastic Remainder Selection, the fitness of an individual f_i is divided by the average fitness of the population \bar{f} . For each string i where f_i/\bar{f} is greater than 1.0, the integer part of the number defines the number of copies of the individual are put forward for crossover. For example an individual with a fitness value of 1.45, places one copy forward for crossover and then has a 0.45 chance of putting a second copy forward.

Stochastic Universal Selection on the other hand, is optimal with respect to efficiency, bias (that is, the distance from the RWS in relation to expected value) and the spread (range of possible individuals put forward for crossover) [9]. SUS simulates a roulette wheel similar to RWS. However, while RWS spins the wheel n times ($n =$ population size), SUS spins the wheel once, using n uniformly spaced pointers at the edge of the wheel. SUS has zero bias, is very efficient and minimises the spread and is regarded as the lowest noise selection scheme.

Figure 8.19 indicates that the SGA (with SRS and SUS individually included), located the local optimum in the initial population, (see the off-line performances) due in part to the level of diversity (illustrated by the SGA on-line performance). However, due to the stochastic nature of the selection mechanisms, the search gets drawn towards the deceptive attractor and as the population converges prematurely




Figure 8.19: Selection Noise On-line/Off-line

as the search continues and diversity is quickly eliminated from the population. Once the landscape changes, because the population, for both the SRS and SUS SGAs, were trapped on the deceptive local optimum, they automatically located to new global optimum and remained there for the duration of the search. The MGA on the other hand located the global optimum both before and after the landscape change.

Selection P-Value Results			
Off-Line	SRS	SUS	MGA
SRS	–	0.04131	0.09138
SUS	0.04131	–	0.4005
On-Line			
SRS	–	0.2732	$< 2.2e^{-16}$
SUS	0.2732	–	$< 2.2e^{-16}$

Table 8.1: Off-Line & On-Line Selection p-values

The statistical analysis of the results between the performances of the MGA and the SGA are shown in Table 8.1 and indicate that the off-line and on-line performances of SRS and SUS are quite similar to one-another. Comparing the MGA’s off-line

results to those of the SRS SGA, shows a less significant result than that of the off-line MGA and SUS SGA. The on-line performance of the MGA differs significantly from both of the SGAs. Figure 8.20, illustrates the rate of convergence within the population and shows both the SGA using SRS and the SGA using SUS, converging quickly. The MGA maintains a higher degree of diversity, both at a phenotypic and genotypic level which assists in successfully locating the global optimum, both before and after the landscape change, as the level of diversity maintained allows the search to escape from the local optimum.




Figure 8.20: Selection Noise Convergence Rate

8.5.5 Selection Pressure

In a small population GA, under normal circumstances it is common to have a number of extraordinary individuals in a population made up primarily of mediocre individuals. With selection, the extraordinary individuals will begin to dominate the population quite quickly and premature convergence will begin to take hold. Even if there is significant diversity within a population, late in a run, the population's

average fitness (on-line) may be close to the population's best fitness (off-line). This leads to a situation where the search for improvement becomes a *random-walk among the mediocre* [67].

Scaling Techniques

Fitness scaling has been used to overcome this problem. We now examine fitness scaling diversifying techniques aimed at reducing the selection pressure within the population:

- Linear scaling adjusts the fitness values of all individuals within the population, such that the fittest individual receives a fixed number of expected offspring and therefore prevents it from reproducing too frequently. The fitness function $f'_i = axf_i + b$ is used where a and b are normally selected so to allow the average individual receive, on average, one offspring copy, and the best receives the specified number of copies (normally two). This method may return a negative fitness value.
- Window scaling, fitness is scaled by subtracting from the raw fitness, the lowest fitness of any individual in the past number of scaling window generations. The fitness function is $f'_i = f_i - f_w$, where w is the window size and is typically somewhere between 2 and 10 and f_w is the worst value observed in the w most recent generations.
- Sigma Truncation avoids returning negative fitness values for individuals within the population and incorporates problem dependant information into the scaling mechanism. The fitness of an individual f'_i is calculated as follows: $f'_i =$

$f_i - (\bar{f} - cx\sigma)$, where c is a small integer value between 1 and 5, \bar{f} is the average raw fitness and σ is the population standard deviation. Negative values for f are avoided as any result $f < 0$ is set to zero. Individuals where $f_i < c$ standard deviation from the average fitness value are not selected.

- Boltzmann tournament selection procedure is derived and implemented to give stable distributions within a population. It also creates another niching mechanism for forming and sizing stable subpopulations of individuals according to differences among them, if the cooling process is not taken to the limit. Boltzmann scaling is expressed as $f' = e^{f_k/T}$ and selection pressure is low when the control parameter T is high.




Figure 8.21: Selection Pressure On-line/Off-line

The results of the scaling experiments are illustrated in Figure 8.21 and Table 8.2. These results indicate that the SGA using various scaling techniques, fails to maintain enough diversity within the population to avoid premature convergence. The statistical results indicate that differences between the various scaling methods,

for these experiments, are minimal. The most significant results are between the MGA and the SGA using each of the scaling mechanisms.

Scaling Comparison					
Off-Line	Linear	Window	Sigma T.	Boltzmann	MGA
Linear	—	1	0.5807	0.5716	$< 2.2e^{-16}$
Window	1	—	0.5807	0.5716	$< 2.2e^{-16}$
Sigma T.	0.5807	0.5807	—	1	$< 2.2e^{-16}$
Boltzmann	0.5716	0.5716	1	—	$< 2.2e^{-16}$
On-Line	Linear	Window	Sigma T.	Boltzmann	MGA
Linear	—	0.9319	0.1071	0.6711	$< 2.2e^{-16}$
Window	0.9319	—	0.04809	0.9433	$< 2.2e^{-16}$
Sigma T.	0.1071	0.04809	—	0.1353	$< 2.2e^{-16}$
Boltzmann	0.6711	0.9433	0.1353	—	$< 2.2e^{-16}$

Table 8.2: Off-Line & On-Line Scaling p-values

Figure 8.22, gives an overview of the rate of convergence associated with the SGA (incorporating scaling mechanisms) and the MGA.




Figure 8.22: Selection Pressure Convergence Rate

The graph illustrates that Linear, Window, Sigma Truncation and Boltzmann SGAs lose diversity very quickly in the search, which results in the failure of the

SGA to adapt after the landscape changes. The MGA implicitly maintains a level of diversity within the population which assists in adapting and locating the global optimum before and after the landscape change.

Niching Techniques

Niching can be defined as an approach which encourages a number of distinct groups of genotypes to develop and remain in the population, with reduced pressure from the GA to converge towards a single type of genotype. Crowding involves a form of niching of the population. The theory behind Crowding is that offspring will have a tendency to replace individuals which are similar [128]. By doing this more often individuals of a similar genotype arise in the population, which increases the chances that their offspring will replace one of them rather than a dissimilar individual. With Crowding, as implemented in this dissertation, before crossover or mutation, normal fitness weighted selection is used to select members for the next generation. Crossover then takes place on individuals selected randomly from this set.

After individuals are selected for crossover, the offspring are calculated as usual. For each offspring, crowding factor, in this case 2, members of the survivors are selected randomly and the Hamming distance of each genotype from the offspring is calculated for 2 individuals. The offspring the replaces whichever survivor is nearest in Hamming distance. Incest reduction is used in conjunction with Crowding and introduces a mechanism to reduce the percentage of crossover between similar genotypes. After being selected, pairs are then selected for crossover by choosing the first parent at random from the list of selected individuals, then choosing a pre-defined

(incest-reduction) number of possible candidates for the other parent randomly. The Hamming distance of each candidate from the first parent and the one with the greatest Hamming distance is selected for crossover.




Figure 8.23: Niching On-line/Off-line

The results of the SGA performance with Crowding and Incest reduction are shown in Figure 8.23. The results indicate that the performance of the SGA with Crowding are very similar to those of the SGA incorporating Incest reduction (off-line p -value = 0.03351 and the on-line p -value = 0.1663). Both niching techniques fail to escape the local optimum. The MGA results differ, in that the diversity maintained within the population allows the search to escape the deceptive trap. Comparing the MGA off-line and on-line results statistically against both of the niching techniques, indicated a p -value $< 2.2e^{-16}$. In relation to the rate of convergence. Unlike the MGA, the niching techniques outlined, lose diversity early in the search (see Figure 8.24), making it difficult for the SGA to adapt and escape the local optimum when the environment changed.




Figure 8.24: Niching Techniques Convergence Rate

8.5.6 Operator Disruption

For the operator disruption experiments, we increased the rate of P_c and P_m (see Figures 8.25 and 8.26 respectively). Examining the results of increasing the rate of crossover to 0.90, whilst maintaining mutation at $1/l$, the SGA located the global optimum early in the search but failed to adapt once the landscape changed.




Figure 8.25: Operator Disruption (Crossover) On-line/Off-line

The results of increasing the rate of mutation to 0.50, while maintaining crossover




Figure 8.26: Operator Disruption (Mutation) On-line/Off-line

at a rate of 0.7, illustrate that the search is being directed by randomness, with the SGA constantly moving to and from the global optimum, shown by the off-line performance. The result also shown a very large degree of variation in the on-line performance. These results appear to indicate that there is too much exploitation in the search for this particular landscape.

The increased crossover results, both off-line and on-line were significantly different from the MGA results, off-line p -value $< 2.2e^{-16}$ and on-line p -value = 0.0003238. One possible reason for this is that higher rates of crossover for this particular problem, result in the balance between exploration and exploitation being negatively effected. The Mutation results differed in comparison with an off-line p -value Of 0.0003238 and on-line p -value $< 2.2e^{-16}$. The mutation off-line results indicate that increased mutation rates are statistically closer to the MGA results, however as the on-line results indicate, the search was driven by randomness and failed to produce good strings, meaning the diversity maintained wasn't useful in terms of a search algorithm.




Figure 8.27: Operator Disruption Convergence Rate

Regarding to the maintenance of variation within the population, Figure 8.27 outlines the impact of operator disruption. Increasing crossover failed to maintain diversity within the population, which converged quickly. The mutation experiment increased the level of diversity within the population, but as mentioned above, the search was directed by randomness. This random diversity failed to assist in the search, unlike the diversity maintained by the MGA.

8.6 Chapter Summary

The aim of this chapter was to examine the population trajectory and the impact on variation and variability of the proposed neutrality in the primary structure of a GA, it also examines the effect on heterogeneity of the proposed representation. The chapter outlined Test 6, which graphically illustrated the population trajectory and showed visually the variation and variability associated with the inclusion of a modular, fixed non-trivial mapping. The experiments in Test 6 were designed to test

the hypothesis: H_{10} *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, alters the evolutionary trajectory and is beneficial over a fully deceptive changing landscape.* The chapter also discussed Test 7, examining heterogeneity and visually represents the evolution of neutral networks, comparing the variation created by the neutral mapping with traditional diversification mechanisms. Test 7 was created to test the hypothesis: H_6 *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, maintains heterogeneity and is beneficial in promoting exploration over a fully deceptive changing landscape.*

The Test 6 experiments examined the impact of the proposed neutrality on population evolution, variation and variability. The test showed visually, the impact of the representation on variation and variability at various stages of the evolutionary trajectory, with results indicating that over the changing deceptive landscape the MGA was able to escape being trapped on the local optima, through it's ability to maintain variation and it's increased phenotypic variability. The results also showed graphically the impact on genotypic variation and phenotypic variation, indicating how the MGA maintains a greater degree of genotypic variation and to a lesser extent an increase in phenotypic variation. The tests compared graphically, the connectivity of the SGA and the MGA, and showed the increased connectivity associated with Missense mutation, illustrating the increase in phenotypic variability. The results obtained from Test 6 indicated that the proposed representation, when compared with that of a SGA, increased variation and variability, which altered the evolutionary trajectory which proved beneficial over a fully deceptive landscape. The findings

achieved show that the stated hypothesis H_{10} is supported.

Test 7 outlined the various neutral networks present in the MGA population and also showed the genotypic neighbourhoods, showing that the neighbourhoods are not randomly created and maintained good solutions within the population. Test 7 examined the impact of neutrality on population heterogeneity by comparing the MGA to a series of SGA's each with a different mechanism to prevent the premature loss of good solutions and sub-solutions. The approaches used by the SGA included various selection mechanisms; selection noise approaches; selection pressure techniques and operator disruption. The selection mechanisms experiments included: Roulette Wheel selection; Tournament selection and Ranking selection, with results showing the variation within the MGA population assists in searching deceptive changing landscapes. The Selection noise experiments included Stochastic Remainder Selection, Stochastic Universal Selection, with results indicating the variation maintained by the MGA outperforming the traditional selection mechanisms used by the SGA over deceptive changing landscapes. Finally, the selection pressure experiments included various scaling techniques and a number of Niching techniques. Results indicated that the Heterogeneity created by the MGA representation outperformed the scaling and niching techniques over the fully deceptive changing landscape. The results obtained from Test 7 indicate that the stated hypothesis H_6 is supported.

From the results presented, neutrality, while increasing the genotype space without increasing the phenotype space, offers a number of potentially useful characteristics for GAs. Results indicate that by incorporation the proposed neutrality, genetic drift slows allowing a GA to maintain additional, highly fit, genetic material within the

population, as the mutation operator produces both silent and adaptive mutations, thereby creating neutral networks of different, but related genotypes. Through an implementation of Neutral theory the MGA implicitly maintains genetic variation within the population through selection and the slowing of genetic drift, which impacts on gene flow and provides additional searching capacity by increasing variability. The proposed MGA mapping provides effective pressure which maintains useful diversity within the population, enabling crossover to combine this diversity with other building blocks.

The findings presented from both Test 6 and Test 7 show that the neutral networks are constructed by combining identical phenotypic structures as reported by Reidys et al. [150] indicating that the most prominent neutral networks have a high degree of connectivity. The results concur also with Shipman [167], who reported that neutrality is advantageous where neutral networks are distributed over the search space with a high degree of connectivity between them. Findings also compare with Nimwegen et al. [143] who showed that by evolving a population through neutral networks, migrating individuals remain highly connected, resulting in phenotypes which are quite resilient against mutation.

These findings suggest a similarity to work by Correia [35], who argues that there are synonymously redundant representations that allow connectivity to be increased between phenotypes, when compared to non-redundant representations. Overall the results presented for Test 6 indicate that by changing the primary structure, the MGA possesses the ability to increase both variation and variability, as a direct result of the presence of neutrality in the genotype-phenotype map, due to the inclusion of

an implementation of transcription, translation and dynamic activation. The results show the influence of a non-trivial neutral representation, where different genotypes in a neutral set allow for different phenotypic distributions [176, 183].

From the results presented, neutrality, while increasing the genotype space without increasing the phenotype space, offers a number of potentially useful characteristics for GAs. Results indicate that by incorporation the proposed neutrality, genetic drift slows allowing a GA to maintain additional, highly fit, genetic material within the population, as the mutation operator produces both silent and adaptive mutations, thereby creating neutral networks of different, but related genotypes. Through an implementation of Neutral theory the MGA implicitly maintains genetic variation within the population through selection and genetic drift, which impacts on gene flow and provides additional searching capacity by increasing variability. The next and final chapter, Chapter 9 discusses the overall findings in relation to previous theory and provides a summary of the conclusions.

Chapter 9

Discussion & Conclusion

9.1 Introduction

Evolution can be viewed as operating on the genotype space while exploring the search space and through recombination and mutation, defines the search space's variational topology. When discussing variational topology we need to examine both genotypic variational topology and phenotypic variational topology. Nature uses a complex genotype-phenotype map to advance a relatively simple genotype space variational topology to an extremely complex phenotypic variational topology.

Chapter 1 motivates the dissertation and outlines biological concepts such as Neutral theory, Modularity, Variation and Variability, intermingling these concepts with GAs to identify a series of open research questions:

- How does the introduction of neutrality, through the use of a biologically inspired genotype-phenotype mapping, impact on a population's evolutionary trajectory over static and dynamic fully deceptive landscapes?

- How, over landscapes of varying degrees of hardness, does the inclusion of neutrality into the genotype-phenotype map of a GA influence genotypic and phenotypic variation?
- How does the introduction of a neutral representation into the primary structure of a GA impact on heterogeneity over a fully deceptive changing landscape?
- How does altering the granularity of neutrality, which introduces varying degrees of neutral drift, impact the phenotypic variability of a genetic algorithm with a many-to-one genotype-phenotype map?
- How does the introduction of neutrality, through the use of a biologically inspired genotype-phenotype mapping impact on the evolutionary path and phenotypic variability over a fully deceptive changing landscape?

The open research questions led to the formulation of a number of hypotheses and defined the primary aims of the dissertation. The hypotheses developed were listed under 3 headings, efficacy, variation and variability, with a search strategy being viewed as beneficial if optimisation is achieved in fewer generations. The hypotheses tested in the dissertation were as follows:

- MGA Efficacy Hypotheses
 - *H₁: Including a biologically inspired, multi-layered, many-to-one genotype-phenotype map into a GA, benefits searching a fully deceptive changing landscapes*

- H_2 : *The inclusion of a biologically inspired, modular, fixed non-trivial, multi-layered genotype-phenotype map into a GA, can perform as robustly as a standard GA in optimising De Jong's test suite.*
- MGA Variation Hypotheses
 - H_3 : *Diversity, implicitly maintained by a many-to-one, genotype-phenotype mapping of a GA, implemented by an interpretation of the biological processes of Transcription and Translation, is beneficial in searching noisy and dynamic landscapes.*
 - H_4 : *Introducing neutrality into the representation of a GA, through a modular, fixed non-trivial mapping, is beneficial in searching static landscapes with varying levels of difficulty.*
 - H_5 : *Including a modular, fixed non-trivial genotype-phenotype map which introduces a more flexible phenotypic structure and a higher degree of phenotypic variation through the use of neutrality, benefits the optimisation of solutions over dynamic landscape problems.*
 - H_6 : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, maintains heterogeneity and is beneficial in promoting exploration over a fully deceptive changing landscape.*
- MGA Variability Hypotheses
 - H_7 : *Altering the level of granularity alters the size of the Translation table and impacts on the search over more difficult dynamic landscapes.*

- H_8 : *The benefit of including a Transcription phase in a many-to-one, genotype-phenotype mapping, increases as the level of problem difficulty increases over more difficult dynamic landscapes.*
- H_9 : *An interpretation of Missense mutation, included within the layering of a biologically inspired, multi-layered genotype-phenotype mapping GA, assists searching more difficult dynamic landscapes.*
- H_{10} : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, alters the evolutionary trajectory and is beneficial over a fully deceptive changing landscape.*

Chapter 1 identified the primary aim of this dissertation as the combination of Darwinism and Neutral theory, through the development of a biologically inspired multi-layered genotype-phenotype map, incorporating the principle of modularity, into a GA and to examine the proposed representation's impact on variation and variability. The motivation derived from the idea of creating a representation which was biologically inspired and involved incorporated aspects of Darwinian survival of the fittest, Neutral theory and Modularity into the primary structure of a GA. As concepts observed in nature often prove useful to GAs when implemented at an abstract level, *transcription*, *translation* and *dynamic activation* were used to produce non-trivial neutrality in the representation and a *missense* mutation operator was also introduced, functioning within the layered mapping. The contribution is to develop a novel way to achieve a balance between exploration and exploitation by creating a genotype-phenotype mapping which provides effective pressure to preserve diversity allowing the continued exchange of building blocks, thereby permitting sustained

exploration. A secondary contribution was the development of a variation operator to operate within the layers of the mapping.

By having a many-to-one genotype-phenotype map, the idea of *phenotypically neutral* neighbours [169] was introduced, which potentially, allow a passage through the genotype space without losing fitness when trapped. As the mapping is designed to maintain variation and increase phenotypic variability, the MGA adopts a modular approach by minimising pleiotropic interaction between characters operating on different functions having a separate representation of character complexes, for distinct functions [192]. The MGA representation adopts a haploid primary structure, thereby allowing standard variational operators to function in a problem independent manner and offers the ability to tune the level of neutrality present, thereby altering the variational topology.

Chapter 2 introduces GA fundamentals and describes how GAs work, outlining concepts such as; Schemata, the Schema Theorem, the Building Block hypothesis, search spaces and fitness landscape. The chapter outlines a number of landscapes chosen to evaluate the proposed representation, including static and dynamic problems, unimodal and multimodal problems, deceptive and non-deceptive problems. Chapter 3 discusses neutrality, referencing past literature and examines Neutral theory and neutrality from a biological prospective. The chapter then expands the discussion to the use of synthetic neutrality in artificial systems, focusing on GAs and illustrating two common approaches to introducing neutrality into a GA; fitness landscapes and genotype-phenotype mappings. The introduction of neutrality through genotype-phenotype mapping is the method adopted in this dissertation.

Chapter 4 describes the design of the proposed representation, using worked examples to illustrate the various mappings of the MGA. The chapter outlines the modular design of the representation and explains how a basic interpretation of the principals of transcription and translation are implemented, producing a layered-mapping and describing the implementation of missense mutation and its associated mappings. The change in variability is also illustrated as the representation is designed to increase connectivity. Finally, a worked example of the creation of a 6-3-1 MGA representation is included, followed by an 8-4-1 MGA representation example, illustrating a number of scenarios involving missense mutation.

9.2 Discussion

A number of tests were designed to examine the hypotheses emerging from the research questions. Chapter 5 begins this process by examining the efficacy of the MGA. Chapter 6 continues and explores the variation associated with the MGA over a number of landscapes, while Chapter 7 examines the variability and accessibility of phenotypes from genotypes by altering the arity of the representation. Chapter 8 concludes the experiments by examining the impact of the proposed representation in relation to evolutionary trajectory and heterogeneity.

9.2.1 Efficacy Results

Test 1 & Test 2

Beginning with Chapter 5, Test 1 examines hypothesis H_1 : *Including a biologically inspired, multi-layered, many-to-one genotype-phenotype map into a GA, benefits searching a fully deceptive changing landscapes.* The results of the Test 1 experiments indicate that the introduction of a flexible phenotypic variational topology enhances the ability of a genetic algorithm to search more complex fully deceptive landscapes. By introducing a more flexible genotypic-phenotypic relationship through the use of a series of mappings, loosely based on the biological processes of transcription and translation, the MGA repeatedly optimised both the fixed deceptive landscape and the changing deceptive landscape. The experiments moved from a basic 3-bit fully deceptive problem to a loosely ordered 10 3-bit fully deceptive problem and finally onto a 10 3-bit fully deceptive changing landscape problem.

Overall the results of Test 1 found that there was little benefit, if any, in introducing neutrality for the 3-bit deceptive problem, mainly due to the relative ease of the problem, as reported by Collins [34]. These results are also in line with Beaudoin et al. [18], who found that as the level of difficulty increased so too did the benefit of including neutrality and Doerr et al. [49], who found that neutrality was possibly beneficial over more difficult deceptive landscapes with multiple local optima. The MGA representation appears to be structured in a way that *improves the likelihood of sampling the optimum* [62], combined with the dampening of the destructive effects of mutation, as reported by Wagner [191] and Yu & Miller [207], with results showing that the proposed multi-layered genotype-phenotype map is effective in solving

the deceptive problems presented. Results also indicated that the benefits associated with the multi-layered genotype-phenotype mapping increases as the level of problem difficulty increases. The experiment results and the statistical analysis indicate that the stated hypothesis (H_1) is supported.

The experiments carried out under Test 2 were designed to test the hypothesis H_2 : *The inclusion of a biologically inspired, modular, fixed non-trivial, multi-layered genotype-phenotype map into a GA, can perform as robustly as a standard GA in optimising De Jong's test suite.* Results show that for the characteristics presented by the Sphere function, the Rosenbrock function, the Step function and the Quadratic function, there appears to be little benefit in introducing neutrality. However, this is not the case for the Shekel's Foxholes experiments, where the introduction of neutrality through the genotype-phenotype mapping has been shown to be beneficial. A possible reason for this is that the neutrality introduced through the multi-layered mapping, reduces the impact of operators such as mutation and crossover, slows genetic drift and assists in avoiding getting stuck in a local optima as the search progresses [191, 207].

The inclusion of an adaptation of the biological concepts of transcription and translation into a GA, introduces neutrality into the genotype-phenotype mapping. The results of the Test 2 experiments over the modified De Jong test suite, indicate classes of problems which could possibly benefit from the inclusion of a multi-layered genotype-phenotype map. The results appear to suggest that the problems most likely to benefit would contain a combination of characteristics such as, multi-dimensionality, multi-modality, non-separable, continuous and deterministic. It also

seems, from the results presented that the MGA is robust, which is a similar finding to Wilke et al. [201] who argued that neutrality provides robustness. The results indicate that the stated hypothesis (H_2) is supported and that the MGA is as robust and capable as a SGA over the test suit.

The results of Tests 1 and 2 have established the efficacy of the MGA over the problem landscapes chosen. Tests 1 and 2 were developed to test hypothesis H_1 and H_2 which were created from the research question, *How does the introduction of neutrality, through the use of a biologically inspired genotype-phenotype mapping, impact on a population's evolutionary trajectory over static and dynamic fully deceptive landscapes?* Results from the experiments supported both hypotheses, indicating that neutrality dampens the effect of variation operators and genetic drift alters as neutrality is introduced, impacting on the evolutionary trajectory. Findings from Tests 1 and 2 appear similar to Banzhaf [10] and Elgin [54], where neutrality in the representation assisted in maintaining diversity and random drift caused by neutral variants increases population diversity by expanding the population distribution so as to help find an escape route from local traps. The findings also concur with Ebner et al. [51, 52] who indicated that neutral networks assist in maintaining diversity in the population, which may be advantageous in a changing environment, and Beaudoin [18] who found that by introducing neutrality easier problems became more difficult and more difficult problems became easier. One reason for the easier problems being more difficult relates to the additional computational overhead associated with neutrality [63].

9.2.2 Variation Results

Test 3 & Test 4

As the findings from Tests 1 and 2 indicated that maintaining variation is an important characteristic of the proposed neutral representation, the next set of experiments, Tests 3 and 4 were designed to examine variation in more details. Test 3 is designed to test hypothesis H_3 : *Diversity, implicitly maintained by a many-to-one, genotype-phenotype mapping of a GA, implemented by an interpretation of the biological processes of Transcription and Translation, is beneficial in searching noisy and dynamic landscapes.* The findings of Test 3 illustrated that variation takes place both at a genotypic level and to much lesser extent, at the phenotypic level, with results indicating that as the levels of problem difficulty increase over fully deceptive noisy and changing landscapes, the benefit of including the proposed representation in searching the search space also increase. The results suggest that the implicit maintenance of variation in the population is one of the reasons for the MGA's success, as variation combined with the multi-layered representation assists in inducing phenotypic variability.

Test 3 experiment results indicated that as problem difficulty increased, so too did the benefit of including neutrality and that easier problems seemed harder and harder problems seemed easier, in a similar manner to Beaudoin [18], with additional computational overhead negatively effecting the easier problems [63]. The results also found that the maintenance of diversity within the population, caused by the dampening of impact from variational operators slowed genetic flow (similar to Yu & Miller [205, 206]) and was beneficial in searching the search space, particularly

over changing landscapes as found by Ebner et al. [51, 52]. The representation presented dictates the effect on genetic flow, which regulates the level of variation and appears to vary according to level of difficulty present in the landscape. In many ways the findings from the experiments are similar to biological studies carried out by King and Dukes [113] who argue that much of the evolution of proteins is down to neutral mutations and genetic drift. The findings suggest that as the level of difficulty increases, the inclusion of neutrality in the representation, by implementing an interpretation of the biological process of transcription and translation, benefits in searching noisy and dynamic landscapes. The Test 3 experiment results support the stated hypothesis (H_3).

Experiments conducted for the first part of Test 4 were created to test hypothesis H_4 : *Introducing neutrality into the representation of a GA, through a modular, fixed non-trivial mapping, is beneficial in searching static landscapes with varying levels of difficulty.* With results indicating that over the static unimodal problem, the neutral mapping was beneficial in searching the landscape as the level of difficulty increased through changes in the dimensionality and that the benefits on including neutrality increased as the level of difficulty increased. The results of the static unimodal experiments illustrated that as problem difficulty increased, the performance of the neutral representation contained in the MGA proved beneficial in searching the search space and illustrated the MGA's robustness as outlined by Wilke et al. [201]. The findings in relation to variation were similar to past research, in that silent mutation is neutral and does not affect the fitness value, but does maintain diversity in the population and this can assist exploration. These findings were similar to those presented by Yu

& Miller [206].

The second part of Test 4 uses a dynamic unimodal problem to test the hypothesis H_5 : *Including a modular, fixed non-trivial genotype-phenotype map which introduces a more flexible phenotypic structure and a higher degree of phenotypic variation through the use of neutrality, benefits the optimization of solutions over dynamic landscape problems.* The results from the second part of Test 4 showed that the dynamic unimodal experiments illustrated the adaptability and robustness of the MGA in the face of a changing landscape, similar to Ebner et al. [51, 52]. The findings highlighted that the benefits associated with the type of neutrality introduced by the layered mapping, increase as the level of difficulty increases and are more pronounced over a changing landscape. Positive results from both parts of Test 4 support both hypotheses (H_4 and H_5).

The hypotheses H_3 , H_4 and H_5 were developed from the research question *How, over landscapes of varying degrees of hardness, does the inclusion of neutrality into the genotype-phenotype map of a GA influence genotypic and phenotypic variation?* Results from the experiments carried out, illustrated that as the level of difficulty increased over the given landscapes the levels of variation maintained assisted in searching the search space. One reason for this is that, the neutral representation, by maintaining variation within the population, primarily at a genotypic level and to a lesser extent at a phenotypic level, allows neutral networks to hold a number of closely related genotypes which, through the layered-mapping, help induce phenotypic variability. The proposed representation creates neutrality through implementations of the biological processes of transcription and translation. These processes had been

applied in artificial systems (i.e. Ashlock [5] and Banzhaf [10]), however the proposed mapping differed significantly as modularity is employed in the way the layers are implemented.

Test 7

Another area related to variability is heterogeneity and is the subject of Test 7, which was created to test the hypothesis H_6 : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, maintains heterogeneity and is beneficial in avoiding premature convergence over a fully deceptive changing landscape.* Test 7 outlined the various neutral networks present in the MGA population and also showed the genotypic neighbourhoods and their proximity to one another in terms of fitness plateaus. Test 7 examined the impact of neutrality on population heterogeneity by comparing the MGA to a series of SGA's each with a different mechanism to prevent the premature loss of good solutions and sub-solutions.

Examining the search space as argued by Toussaint [183] allows an insight to be developed into how evolution can adapt the search on the phenotypic search space and results indicate that the variation in the population contains genotypes which are close to one another in terms of the phenotypes they represent rather than their genetic makeup. The findings show that the neutral networks are constructed by combining identical phenotypic structures as reported by Reidys et al. [150] and illustrate that the most prominent neutral networks have a high degree of connectivity. The results also concur with findings by Shipman [167], who argued that neutrality

is advantageous where neutral networks are distributed over the search space with a high degree of connectivity between them. The examination of the population's evolutionary trajectory from the phenotype space illustrates that the MGA's representation maintains good solutions and sub-solutions within the population allowing the processing of good schemata, remembering also the Building Block hypothesis, which states that the power of a GA lies in being able to find good *building blocks* [67], which can be used for continued exploration.

The various diversifying techniques used by the SGA in Test 7, included a number of selection mechanisms; selection noise approaches; selection pressure techniques and operator disruption. The selection mechanisms experiments included: Roulette Wheel selection; Tournament selection and Ranking selection, with results showing the variation within the MGA's population assists in searching deceptive changing landscapes. The Selection noise experiments included Stochastic Remainder Selection, Stochastic Universal Selection, with results indicating the variation maintained by the MGA outperforming the traditional selection mechanisms used by the SGA over deceptive changing landscapes. The results in relation to the neutrality maintenance of diversity are similar to those reported by Yu and Miller [205, 206] and Ebner et al. [51, 52] in relation to neutrality and changing landscapes.

Finally, the selection pressure experiments included various scaling techniques and a number of Niching techniques. Again, results indicated that the heterogeneity created by the genotype-phenotype mapping outperformed the scaling and niching techniques over the fully deceptive changing landscape. Overall, the heterogeneity experiment results illustrate that the multi-layered mapping, through neutrality, im-

pacts in a positive way on heterogeneity by maintaining variation within the population, with the neutral network results showing that the representation affects genetic drift while still promoting highly fit schema in the population. The hypothesis (H_6) emerged from the research question *How does the introduction of a neutral representation into the primary structure of a GA impact on heterogeneity over a fully deceptive changing landscape?* and results of the experiments carried out in Test 7 resulted in the stated hypothesis (H_6) being supported.

9.2.3 Variability Results

Test 5

Test 5 looks at the impact of altering the arity of the MGA, which alters the granularity of the representation, changing the accessibility of phenotypes from genotypes and tests hypothesis H_7 : *Altering the level of granularity alters the size of the Translation table and impacts on the search over more difficult dynamic landscapes.* Test 5 also examines the transcription phase of the mapping and test hypothesis H_8 : *The benefit of including a Transcription phase in a many-to-one, genotype-phenotype mapping, increases as the level of problem difficulty increases over more difficult dynamic landscapes.* Finally, Test 5 examines the impact of the missense mutation operator and tests the hypothesis H_9 : *An interpretation of Missense mutation, included within the layering of a biologically inspired, multi-layered genotype-phenotype mapping GA, assists searching more difficult dynamic landscapes.*

The results of Test 5 indicate that overall the MGA appears robust over both static and changing landscapes. By incorporating a modular, fixed non-trivial tun-

able genotype-phenotype mapping, the MGA offers the ability to tune the granularity of the representation, which appears beneficial over various landscapes and the performance of the MGA varies, depending on the level of granularity in the representation and the problem difficulty. Over the static OneMax landscapes, the MGA's off-line and on-line performances were quite similar to those of the SGA. One reason for this is that the unimodal OneMax problem, by its nature, is relatively easy for a GA to solve.

The results illustrated that there was little benefit in including neutrality on unimodal landscapes and that easier unimodal problems appeared more difficult and more difficult unimodal problems became easier, this trend continued as arity increased. These results are consistent with those found by Beaudoin et al. [18] and that little exploration was required due to the level of problem difficulty associated with the unimodal landscapes. The results also illustrated that the MGA's representation reduces the impact of variational operators, as reported by Doerr et al. [49]. Also, as found by Galván-López and Poli [63] the increase in computational overhead associated with the introduction of neutrality out-weights the introduction of neutrality on static unimodal landscapes. The results indicated that as the level of problem difficulty increased over the static unimodal landscape, the number of neutral networks decreased, which ultimately results in a reduction in overall variability.

However, over the loosely ordered 10 3-bit deceptive problem and particularly over the more difficult 30 3-bit deceptive problem, the MGA out-performed the SGA. It should be noted that the 4-2-1 MGA representation which struggled slightly after the landscape changed, indicating that the adaptability of the MGA decreases when too

coarse a granularity is adopted. It is interesting to note that the MGA, in terms of robustness, performed as well as the SGA over the OneMax landscapes. However, over the deceptive changing landscape experiments the MGA outperformed the SGA, both on-line and off-line, particularly over the more difficult 30 3-bit deceptive problem. Overall, with the more challenging dynamic multimodal deceptive problems, the MGA outperformed the SGA and the performance improved as the level of problem difficulty increased. It appears that given the level of difficulty associated with the landscape, the addition of neutrality was beneficial, which again is similar to findings by Beaudoin et al. [18]. The results over the multimodal deceptive landscape indicated that as arity increased, the degree of exploration required decreased over the less difficult problem and increased over the more difficult problem.

These results suggest that the impact of arity differs depending on the level of problem difficulty. They also suggest that arity impacts on neutral networks, with the number of neutral networks present increasing as arity increases over more difficult problems, resulting in increased exploration. These findings may be explained by biological studies carried out by Huynen [100] which showed with large amounts of neutrality present in the mapping, neutral paths exist, which allow for smooth exploration.

The hypotheses tested (H_7 , H_8 & H_9) emerged for the research question *How does altering the granularity of neutrality, which introduces varying degrees of neutral drift, impact the phenotypic variability of a genetic algorithm with a many-to-one genotype-phenotype map?*. In relation to altering the size of the translation table, results indicate that there is an impact on search particularly over the more difficult

landscapes, so the hypothesis (H_7) can be accepted. The experiments on the transcription phase produced interesting results in that over the easier 10 3-bit deceptive problem, the exclusion of a transcription phase made little or no difference. However over the more difficult 30 3-bit deceptive problem the MGA struggled without the transcription phase, suggesting that the mapping provided by transcription benefits the search strategy, possibly in that it replicates a form of inversion. The results indicated that the hypothesis (H_8) is accepted. Finally, the missense mutation experiments indicated that its presence assisted the search strategy as the optimum was located with fewer functional evaluations. Therefore, when it was omitted over the duration of the run the search failed, allowing the hypothesis (H_9) to be supported.

Test 6

Test 6 was created to graphically illustrate the impact of multi-layered neutral mapping on population trajectory, variation and variability and tests the hypothesis H_{10} : *A modular, fixed non-trivial genotype-phenotype mapping, which introduces neutrality into the primary structure of a GA, alters the evolutionary trajectory and is beneficial over a fully deceptive changing landscape.* The hypothesis tested H_{10} came from the research question *How does the introduction of neutrality, through the use of a biologically inspired genotype-phenotype mapping impact on the evolutionary path and phenotypic variability over a fully deceptive changing landscape?* The results showed visually, the impact of the representation on variation and variability at various stages of the evolutionary trajectory, indicating that over the changing deceptive landscape the MGA avoided being trapped on the local optima, through

it's ability to maintain variation and its increased phenotypic variability. This work compares favourably with Nimwegen et al. [143] who found that during the evolution of a population through neutral networks, migrating individuals remain highly connected, resulting in phenotypes which are quite resilient against mutation. Wagner [191], took a similar view and suggested that neutrality assisted in sheltering a system from the impact of mutation and argued that new adaptations, provided by neutrality, can assist the search. The work also compares with that of Wilke et al. [201], who argued that neutrality provides robustness, with results indicating that genotypes closely connected by mutation, have high selection rates and that these genotypes tend to be located on flatter areas of the landscape.

The Test 6 findings show visually an interpretation of Toussaint's [184] idea of non-trivial neutrality, where different genotypes in a neutral set allow for different phenotypic distributions and how evolution can adapt the search on the phenotype space [183]. Although two genotypes can be considered equivalent if they represent the same phenotype, however their variational topologies can differ [176]. But ultimately the genotype-phenotype map *induces a variational topology on the phenotype space depending on the topology of the genotype space* [183]. The results of Test 6 illustrate graphically the variation and variability associated with the MGA. The findings here suggest that through the neutral representation and dynamic activation, the representation which is designed to be uniform and increase connectivity, alters the population distribution due to changes in the evolutionary trajectory, showing that the performance of the GA can differ even when no individuals are over-represented. The results of Test 6 indicate that hypothesis H_{10} is supported.

To summarise the results of Test 6 which showed graphically, the impact on genotypic variation and phenotypic variation, indicating how the MGA maintains a greater degree of genotypic variation and to a lesser extent phenotypic variation. The tests compared graphically, the connectivity of the SGA and the MGA, and showed the increased connectivity associated with the proposed mapping, illustrating the increase in phenotypic variability. The contribution of this form of mapping lies in its ability to maintain variation without introducing a random search strategy and increase variability, thereby allowing the occupation by the population, of a greater number of fitness plateaus and preventing premature convergence through changing the genetic flow. By adopting this approach, convergence at a phenotypic level can be achieved, but genetic diversity is maintained at a genotypic level. Neutral theory [111], would suggest that where genetic changes spread across a population, changes may or may not alter the phenotype and are a result of genetic drift.

9.3 Conclusion

Given that a GA search involves a mapping between the genotype and the phenotype, a SGA, because of its one-to-one genotype-phenotype mapping, quickly eliminates diversity from the population through its selection policy and low mutation rates. However, variation within the population is a desirable feature as a population's ability to survive often depends on a level of diversity maintained within the population. Therefore, as phenotypic variation is critical for evolution and the genetic representation of a trait determines the variability of a phenotype rather than the genetic variation within the population, then because both variation and vari-

ability are influenced by the genetic representation, the choice of representation is a critical component of the search strategy.

The proposed redundant representation adopted a modular approach which included little pleiotropy, with a separate genetic representation of characters for each phenone so as to improve evolvability through limiting the interference between adaptation of characters [192]. By including Modularity and Neutral theory into the genetic representation the aim is to achieve a balance between exploration and exploitation by increasing variation and variability within the population, thereby improving the adaptability of the algorithm. The representation developed incorporates Neutral theory into a tunable genotype-phenotype mapping through the biological concepts of *transcription* and *translation*.

The results presented, illustrate that through the implementation of Neutral theory, as proposed by Kimura [111], the multi-layered genotype-phenotype mapping presented allows for a tunable, modular, fixed non-trivial relationship. The incorporation of Neutral theory into the MGA genotype-phenotype mapping creates a representation, which increases connectivity and allows variational operators function in a problem independent way, maintaining highly fit schemata within the population. The layered neutral representation, dampens the effect of the variation operators and alters genetic drift, impacting on variation and variability. Neutral theory suggests that as changes at a genotypic level may or may not alter the phenotype, neutral networks emerge. It is the presence of these neutral networks and the altering of genetic drift which exerts a great influence on the evolutionary trajectory, thereby accounting for the MGA's ability to succeed over more difficult and changing landscapes.

The primary aim of this dissertation was to combine Darwinian survival of the fittest [39] with Neutral theory [111] in a tunable, modular, fixed non-trivial GA representation, through a basic implementation of the biological processes of translation and transcription, which allows genetic drift to alter without developing a random search strategy. The representation was designed to allow operators function correctly through uniformity and synonymity, and also to include a high degree of locality and connectivity.

A secondary aim was to introduce variation within the layered architecture through a novel implementation of missense mutation and to examine its impact. The proposed GA representation included neutrality which allows the mutation operator to produce both silent and adaptive mutations, thereby creating neutral networks of different but related genotypes, which through an interpretation of biological concepts, maintained variation and increased variability over the chosen test suites. The mapping benefitted the search strategy over the more difficult changing landscapes, with results indicating that the proposed representation altered the evolutionary trajectory of the population. In general the proposed representation delayed to convergence of the strategy, altered the evolutionary trajectory and changed the long term behaviour of the algorithm. Findings presented showed the benefit to the search strategy of including a variation operator within the layers of the mapping, with missense mutation increasing variability, without drifting towards randomness.

To conclude, through the proposed combination of Darwinism and Neutral theory, the MGA, through a biologically inspired multi-layered representation, produced an effective pressure to maintain useful diversity within the population, allowing that

diversity to be crossed with other building blocks, thereby permitting sustained exploration. The neutral representation created neutral networks and altered genetic drift, which impacted on gene flow and changed the evolutionary trajectory by inducing variability.

9.4 Future Work

Further work emerging from this dissertation includes the following:

- Detailed examination into the self-adaptive nature of a neutral fixed non-trivial representation and the introduction of self-adaptive operators within the proposed MGA. Much research has been carried out into the use of self adaption in EA's, however as pointed out by Toussaint and Igel [185], approaches to self-adaption in evolutionary algorithms can be viewed as an example of the benefits of neutrality. As chromosome can have an associated mutation rate and may be part of the same neutral network, as a result of the presence of neutrality. This results in a situation where chromosomes in the same neutral network may have and have different mutation rates and evolution can choose between these in a self-adaptive way. From a research perspective this is interesting as there exists a variety of mutation rates or distributions within the population, which can evolve.
- Further exploration into the use of *Dynamic Activation* within EAs, which allows the possibility of loci on a genotype, where given a certain context are *functionless (mutation makes no difference)* [84], but combined with a value

elsewhere on the genotype may become important [84]. Combined with this would be an examination into the efficacy of a basic interpretation of the biological concepts of *exons* and *introns* (Exons can be described as coding regions which are interpreted by non coding regions, known as introns), and the development and examination of a mutation operator based on the biological *nonsense mutation* operator, which creates a nonsense codon that does not code for an amino acid.

- As the genotype-phenotype mapping is central to the representation problem and determines the evolvability of the phenotype [192], another avenue for research involves the use of numerical values in MGA's layered representation rather than characters and an examination of the impact on variation and variability. This offers a novel approach to the translation phase which impacts on the accessibility of phenotypes from genotypes and offers the possibility of further insight into the use of both the transcription process and the translation process.

Bibliography

- [1] ALBERCH, P. From genes to phenotype: dynamical systems and evolvability. *Genetica* 84, 1 (1991), 5–11.
- [2] ALTENBERG, L. NK Fitness Landscapes. In *The Handbook of Evolutionary Computation, Section B.2.7.2*, e. a. T. Bäck, Ed. Oxford University Press, 1997.
- [3] ANGELINE, P. J. Adaptive and self-adaptive evolutionary computations. In *Computational Intelligence: A Dynamic Systems Perspective*, M. Palaniswami and Y. Attikiouzel, Eds. IEEE Press, 1995, pp. 152–163.
- [4] ANGELINE, P. J., FOGEL, D. B., AND FOGEL, L. J. A comparison of self-adaptation methods for finite state machines in dynamic environments. In *in Proc. 5th Annu. Conf. Evolutionary Programming* (Cambridge, MA, 1996), L. J. Fogel, P. J. Angeline, and T. Bäck, Eds., MIT Press.
- [5] ASHLOCK, D., SCHONFELD, J., AND MCNICHOLAS, P. D. Translation tables: A genetic code in a evolutionary algorithm. In *IEEE Congress on Evolutionary Computation* (2011), pp. 2685–2692.
- [6] BÄCK, T. Optimal mutation rates in genetic search. In *Proceedings of the First European Conference on Artificial Life* (Cambridge, MA, 1992), F. Varela and P. Bourguine, Eds., MIT Press, pp. 2–8.
- [7] BÄCK, T. Self-adaption in genetic algorithms. In *Proc. of the Fifth Int. Conf. on Genetic Algorithms* (San Mateo, CA, 1993), S. Forrest, Ed., Morgan Kaufmann, pp. 2–8.
- [8] BÄCK, T., FOGEL, D. B., AND MICHALEWICZ, Z., Eds. *Handbook of Evolutionary Computation*, 1st ed. IOP Publishing Ltd., Bristol, UK, UK, 1997.
- [9] BAKER, J. E. Adaptive Selection Methods for Genetic Algorithms. In *Proc. of the International Conference on Genetic Algorithms and Their Applications* (Pittsburgh, PA, 1985), pp. 101–111.

-
- [10] BANZHAF, W. Genotype-phenotype-mapping and neutral variation: A case study in genetic programming. In *PPSN III: Proceedings of the 3rd International Conference on Parallel Problem Solving from Nature* (Jerusalem, Israel, 1994), Y. Davidor, H. Schwefel, and R. Manner, Eds., Springer-Verlag, pp. 322–332.
- [11] BANZHAF, W., AND LEIER, A. Evolution on neutral networks in genetic programming. In *Genetic Programming Theory and Practice III* (Boston, MA, 2006), T. Yu, R. Riolo, and B. Worzel, Eds., Springer US, pp. 207–221.
- [12] BARKER, A. L., AND MARTIN, W. N. Dynamics of a distance-based population diversity measure. In *Proceedings of the 2000 Congress on Evolutionary Computation* (Piscataway, NJ, 2000), IEEE Service Centre, pp. 1002–1009.
- [13] BARNETT, L. Ruggedness and Neutrality – the NKp family of Fitness Landscapes. In *Artificial Life VI: Proc. of the Sixth Int. Conf. on Artificial Life* (Cambridge, MA, 1998), C. Adami, R. K. Belew, H. Kitano, and C. Taylor, Eds., The MIT Press, pp. 18–27.
- [14] BARNETT, L. Netcrawling-optimal evolutionary search with neutral networks. In *In Proceedings of the 2001 Congress on Evolutionary Computation* (2001), IEEE Press, pp. 30–37.
- [15] BARRY, P. Genome 2.0: Mountains of new data are changing old views. *Science News* 172, 10 (2007), 154.
- [16] BEASLEY, D., BULL, D., AND MARTIN, R. An overview of genetic algorithms part 1 fundamentals. Tech. rep., University of Purdue, 1993.
- [17] BEASLEY, D., BULL, D., AND MARTIN, R. An overview of genetic algorithms part 2 research topics. Tech. rep., University of Purdue, 1993.
- [18] BEAUDOIN, W., VEREL, S., COLLARD, P., AND ESCAZUT, C. Deceptiveness and neutrality the nd family of fitness landscapes. In *GECCO'06* (2006), pp. 507–514.
- [19] BENDTSEN, C. N., AND KRINK, T. Dynamic memory model for non-stationary optimization. In *In Proceedings of the 2002 Congress on Evolutionary Computation* (2002), pp. 145 – 150.
- [20] BETHKE, A. D. *Genetic algorithms as functional optimisers*. PhD thesis, (Doctoral dissertation, University of Michigan) Dissertation Abstracts International, 41(9),3503B. (University Microfilms No. 81-06101), 1981.
- [21] BLICKLE, T., AND THIELE, L. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation* 4, 4 (1997), 361–394.

- [22] BRANKE, J. Memory enhanced evolutionary algorithms for changing optimization problems. In *Proceedings of the Congress on Evolutionary Computation* (Washington D.C., USA, 1999), P. J. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, and A. Zalzal, Eds., vol. 3, IEEE Press, pp. 1875 – 1882.
- [23] BRANKE, J. Evolutionary approaches to dynamic optimization problems - updated survey. In *GECCO Workshop on Evolutionary Algorithms for Dynamic Optimization Problems* (01 2001).
- [24] BRANKE, J., KAUSSLER, T., SCHMIDT, C., AND SCHMECK, H. A multi-population approach to dynamic optimization problems. In *In Proceedings of the 4th International Conference on Adaptive Computing in Design and Manufacturing* (2000), pp. 299 – 308.
- [25] BRANKE, J., SALIHOĞLU, E., AND UYAR, C. Towards an analysis of dynamic environments. In *Proceedings of the 7th Annual Conference on Genetic and Evolutionary Computation* (New York, NY, USA, 2005), GECCO '05, ACM, pp. 1433–1440.
- [26] BURKE, D. S., DE JONG, K. A., GREFENSTETTE, J. J., RAMSEY, C. L., AND WU, A. S. Putting more genetics into genetic algorithms. *Evol. Comput.* 6, 4 (1998), 387–410.
- [27] CANGELOSI, A., PARISI, D., AND NOLFI, S. Cell division and migration in a 'genotype' for neural networks (cell division and migration in neural networks). *Network* 5 (1994), 497–515.
- [28] CHEN, S., AND SMITH, F. S. Putting the 'genetics' back into genetic algorithms (reconsidering the role of crossover in hybrid operators). In *Foundations of Genetic Algorithms 5*, W. Banzhaf and C. Reeves, Eds. Morgan Kaufmann., 1999.
- [29] CHOW, R. Effects of phenotypic feedback and the coupling of genotypic and phenotypic spaces in genetic searches. In *Proceedings of the 2004 IEEE Congress on Evolutionary Computation (CEC-2004)* (Portland, Oregon, USA, 2004), vol. 1, IEEE, pp. 242–249.
- [30] CHOW, R. Evolving genotype to phenotype mappings with a multiple-chromosome genetic algorithm. In *Genetic and Evolutionary Computation - GECCO 2004*, vol. 3102 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2004, pp. 1006–1017. 10.1007/978-3-540-24854-5_100.
- [31] COBB, H. G. An investigation into the use of hypermutation as an adaptive operator in genetic algorithms having continuous, time-dependent nonstationary

- environments. Tech. Rep. AIC-90-001, Naval Research Laboratory, Washington, USA, 1990.
- [32] COBB, H. G., AND GREFENSTETTE, J. J. Genetic Algorithms for Tracking Changing Environments. In *Proc. of the 5th Int. Conf. on Genetic Algorithms and their Applications* (1993), Morgan Kaufmann, pp. 523–530.
- [33] COHOON, J. P., HEGDE, S., MARTIN, W. N., AND RICHARDS, D. Distributed genetic algorithms for the floorplan design problem. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 10, 4 (Apr 1991), 483–492.
- [34] COLLINS, M. Finding needles in haystacks is harder with neutrality. *Genetic Programming and Evolvable Machines* 7 (June 2006), 131–144.
- [35] CORREIA, M. B. A study of redundancy and neutrality in evolutionary optimization. *Evolutionary Computation* 21, 3 (Sept 2013), 413–443.
- [36] CORREIA, M. B. A study of redundancy and neutrality in evolutionary optimization. *Evol. Comput.* 21, 3 (2013), 413–443.
- [37] CORREIA, M. B. New insights on neutral binary representations for evolutionary optimization. *Theoretical Computer Science*, 641 (2016), 53–74.
- [38] DA GRACA LOBO, F. M. P. *The Parameter-Less Genetic Algorithm: Rational And Automated Parameter Selection For Simplified Genetic Algorithm Operation*. PhD thesis, University of Illinois at Urbana-Champaign, 2000.
- [39] DARWIN, C. *On the Origin of Species by means of Natural Selection*. John Murray, Albemarle Street, London, 1859.
- [40] DASGUPTA, D. Handling deceptive problems using a different genetic search. In *Proceedings of the First IEEE Conference on Evolutionary Computing vol.1* (1994), pp. 807–811 vol.2.
- [41] DASGUPTA, D. Incorporating redundancy and gene activation mechanism in genetic search for adapting to non-stationary environments. In *Practical Handbook of Genetic Algorithms* (1995), L. Chambers, Ed., CRC Press, pp. 303–316.
- [42] DASGUPTA, D., AND MCGREGOR, D. R. Nonstationary function optimization using the structured genetic algorithm. In *PPSN* (1992), pp. 147–156.
- [43] DE JONG, K. A. *An Analysis of the Behavior of a Class of Genetic Adaptive Systems*. PhD thesis, University of Michigan, Ann Arbor, 1975. Dissertation Abstracts International 36(10), 5140B; UMI 76-9381.

- [44] DE JONG, K. A. Genetic algorithms are NOT function optimizers. In *Foundations of Genetic Algorithms 2* (San Mateo, CA, 1993), L. D. Whitley, Ed., Morgan Kaufmann Publishers.
- [45] DE JONG, K. A. Evolving in a changing world. In *Foundations of Intelligent Systems: 11th International Symposium, ISMIS'99 Warsaw, Poland, June 8–11, 1999 Proceedings* (Berlin, Heidelberg, 1999), Z. W. Raś and A. Skowron, Eds., Springer Berlin Heidelberg, pp. 512–519.
- [46] DE JONG, K. A., POTTER, M. A., AND SPEARS, W. M. Using problem generators to explore the effects of epistasis. In *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)* (San Francisco, CA, 1997), T. Bäck, Ed., Morgan Kaufmann.
- [47] DEB, K., AND GOLDBERG, D. E. Analyzing deception in trap functions. In *FOGA* (1992), pp. 93–108.
- [48] DOERR, B., GAO, W., AND NEUMANN, F. Runtime analysis of evolutionary diversity maximization for oneminmax. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016* (New York, NY, USA, 2016), GECCO '16, ACM, pp. 557–564.
- [49] DOERR, B., GNEWUCH, M., HEBBINGHAUS, N., AND NEUMANN, F. A rigorous view on neutrality. In *IEEE Congress on Evolutionary Computation'07* (2007), pp. 2591–2597.
- [50] E. GOLDBERG, D., DEB, K., AND HORN, J. Massive multimodality, deception, and genetic algorithms. *Parallel Problem Solving from Nature 2* (08 1994), 37–48.
- [51] EBNER, M., LANGGUTH, P., ALBERT, J., SHACKLETON, M., AND SHIPMAN, R. On Neutral Networks and Evolvability. In *IEEE Congress on Evolutionary Computation (CEC)* (2001), IEEE Press.
- [52] EBNER, M., SHACKLETON, M., AND SHIPMAN, R. How neutral networks influence evolvability. *Complex*. 7, 2 (2001), 19–33.
- [53] EIBEN, A. E., HINTERDING, R., AND MICHALEWICZ, Z. Parameter control in evolutionary algorithms. *IEEE Trans. on Evolutionary Computation* 3, 2 (1999), 124–141.
- [54] ELGIN, M. *Steps towards Life: a Perspective on Evolution*. Oxford University Press, 1992.

- [55] ESHELMAN, L. J., AND SCHAFFER, D. J. Preventing premature convergence in genetic algorithms by preventing incest. In *Proceedings of the Fourth International Conference on Genetic Algorithms* (1991), R. K. Belew and L. B. Booker, Eds., San Francisco, CA: Morgan Kaufmann, pp. 115–122.
- [56] FOGEL, D. B. Phenotypes, genotypes, and operators in evolutionary computation. In *Proceedings, IEEE* (1995).
- [57] FONSECA, C. M., AND CORREIA, M. B. Developing redundant binary representations for genetic search. In *In Proceedings of the 2005 IEEE Congress on Evolutionary Computation (CEC 2005)* (Edinburgh, September 2005), IEEE, pp. 372–379.
- [58] FONSECA, C. M., AND CORREIA, M. B. Developing redundant binary representations for genetic search. In *2005 IEEE Congress on Evolutionary Computation* (Sept 2005), vol. 2, pp. 1675–1682 Vol. 2.
- [59] FONTANA, W., AND SCHUSTER, P. Continuity in evolution: On the nature of transitions. In *Science 280:14311452* (1998.).
- [60] FONTANA, W., AND SCHUSTER, P. Shaping Space: The Possible and the Attainable in RNA Genotype-Phenotype Mapping. Working Papers ir98004, International Institute for Applied Systems Analysis, Feb 1998.
- [61] GALVÁN-LÓPEZ, E. Effects of neutrality on evolutionary search. In *In EvoPhD - Second European Graduate Student Workshop on Evolutionary Computation* (Valencia, Spain, April 2007).
- [62] GALVÁN-LÓPEZ, E., AND POLI, R. An Empirical Investigation of How and Why Neutrality Affects Evolutionary Search. In *Proceedings of the 2006 Conference on Genetic and Evolutionary Computation GECCO 2006* (Seattle, WA, USA, 2006), ACM Press, pp. 1149–1156.
- [63] GALVÁN-LÓPEZ, E., AND POLI, R. Some steps towards understanding how neutrality affects evolutionary search. In *Parallel Problem Solving from Nature - PPSN IX*, T. Runarsson, H.-G. Beyer, E. Burke, J. Merelo-Guervs, L. Whitley, and X. Yao, Eds., vol. 4193 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2006, pp. 778–787.
- [64] GEARD, N., WILES, J., HALLINAN, J., TONKES, B., AND SKELLETT., B. A comparison of neutral landscapes nk, nkp and nkq. In *Proceedings of Congress on Evolutionary Computation (CEC 2002)* (2002), vol. 1, IEEE Press, pp. 205–210.
- [65] GOLDBERG, D. E. Simple genetic algorithms and the minimal deceptive problem. *Genetic Algorithms and Simulated Annealing* (1987), 74–88.

-
- [66] GOLDBERG, D. E. Genetic algorithms and walsh functions: Part ii, deception and its analysis. *Complex Systems* 3(2) (1989), 153–171.
- [67] GOLDBERG, D. E. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Inc., Reading, MA, 1989.
- [68] GOLDBERG, D. E. Construction of high-order deceptive functions using low-order walsh coefficients. *Annals of Mathematics and Artificial Intelligence* 5, 1 (Mar 1992), 35–47.
- [69] GOLDBERG, D. E. *The Design of Innovation*. Kluwer Academic Publisher, Boston, 2002.
- [70] GOLDBERG, D. E., AND BRIDGES, C. L. An analysis of a reordering operator on a ga-hard problem. *Biological Cybernetics* 62, 5 (Mar 1990), 397–405.
- [71] GOLDBERG, D. E., AND DEB, K. A comparative analysis of selection schemes used in genetic algorithms. In *Foundations of Genetic Algorithms* (1991), Morgan Kaufmann, pp. 69–93.
- [72] GOLDBERG, D. E., DEB, K., KARGUPTA, H., AND HARIK, G. Rapid, accurate optimization of difficult problems using fast messy genetic algorithms. *Proceedings of the Fifth International Conference on Genetic Algorithms* (1993), 56–64.
- [73] GOLDBERG, D. E., KORB, B., AND DEB, K. Messy genetic algorithms: Motivation, analysis, and first results. *Complex Systems* 3, 5 (1990), 493–530.
- [74] GOLDBERG, D. E., AND RICHARDSON, J. Genetic Algorithms with Sharing for Multimodal Function Optimization. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic algorithms and their application* (Hillsdale, NJ, USA, 1987), L. Erlbaum Associates Inc., pp. 41–49.
- [75] GOLDBERG, D. E., AND SEGREST, P. Finite markov chain analysis of genetic algorithms. In *Proceedings of the Second International Conference on Genetic Algorithms* (1987), pp. 1–8.
- [76] GOLDBERG, D. E., AND SMITH, R. E. Nonstationary Function Optimization Using Genetic Algorithm with Dominance and Diploidy. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application* (Hillsdale, NJ, USA, 1987), L. Erlbaum Associates Inc., pp. 59–68.
- [77] GOULD, S. J. Exaptation: A crucial tool for an evolutionary psychology. *Journal of Social Issues* 47, 3 (1991), 43–65.

- [78] GREFENSTETTE, J. J. Optimization of control parameters for genetic algorithms. *IEEE Transactions on Systems, Man and Cybernetics SMC-16(1)* (1986), 122–128.
- [79] GREFENSTETTE, J. J. Genetic algorithms for changing environments. In *Parallel Problem Solving from Nature 2 (Proc. 2nd Int. Conf. on Parallel Problem Solving from Nature, Brussels 1992)* (Amsterdam, 1992), R. Männer and B. Manderick, Eds., Elsevier, pp. 137–144.
- [80] GREFENSTETTE, J. J. Deception considered harmful. In *Foundations of Genetic Algorithms 2*, L. D. Whitley, Ed. Morgan Kaufmann, San Mateo, CA, 1993, pp. 75–91.
- [81] GREFENSTETTE, J. J., AND BAKER, J. How genetic algorithms work: a critical look at implicit parallelism. In *Proc. of the Third. International. Conf. on Genetic Algorithms* (1989), Morgan Kaufmann.
- [82] GRUAU, F., LYON I, L. C. B., DOCTORAT, O. A. D. D., DEMONGEOT, M. J., COSNARD, E. M. M., MAZOYER, M. J., PERETTO, M. P., AND WHITLEY, M. D. Neural network synthesis using cellular encoding and the genetic algorithm., 1994.
- [83] HAMMING, R. W. Error detecting and error correcting codes. *The Bell System Technical Journal* 29, 2 (April 1950), 147–160.
- [84] HARVEY, I., AND THOMPSON, A. Through the labyrinth evolution finds a way: A silicon ridge. In *Proceedings of the First International Conference on Evolvable Systems: From Biology to Hardware* (1996), vol. 1259, Springer Verlag, pp. 406–422.
- [85] HATZAKIS, I., AND WALLACE, D. Dynamic multi-objective optimization with evolutionary algorithms: A forward-looking approach. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation, Seattle, Washington, USA* (New York, NY, USA, 2006), GECCO '06, ACM, pp. 1201–1208.
- [86] HILL, S., AND O'RIORDAN, C. A genetic algorithm with a multi-layered genotype-phenotype mapping. In *International Conference on Evolutionary Computation (ICEC 2010)* (24-26 Oct 2010), Valencia, Spain.
- [87] HILL, S., AND O'RIORDAN, C. Solving fully deceptive problems in changing environments. In *Artificial Intelligence Cognitive Studies (AICS)* (2010), pp. 87–95.

- [88] HILL, S., AND O’RIORDAN, C. Examining the use of a non-trivial fixed genotype-phenotype mapping in genetic algorithms to induce phenotypic variability over deceptive uncertain landscapes. In *Proceedings of the 2011 Congress of Evolutionary Computation (CEC 2011)* (5-8 June 2011), New Orleans, USA.
- [89] HILL, S., AND O’RIORDAN, C. Neutrality through transcription and translation in genetic algorithm representation. In *4th International Conference on Evolutionary Computation Theory and Applications ECTA Barcelona, Spain* (2012).
- [90] HILL, S., AND O’RIORDAN, C. Analysing the impact of dimensionality and diversity in a multi-layered genotype-phenotype mapped ga. In *CEC* (June 2013), Cancun, Mexico.
- [91] HILL, S., AND O’RIORDAN, C. Altering the granularity of neutrality in a multi-layered genetic algorithm. In *6th International Conference on Evolutionary Computation Theory and Applications ECTA Rome, Italy* (2104).
- [92] HILL, S., AND O’RIORDAN, C. Examining the impact of neutral theory on genetic algorithm population evolution. In *7th International Conference on Evolutionary Computation Theory and Applications ECTA Lisbon, Portugal* (2105).
- [93] HILL, S., AND O’RIORDAN, C. Diversifying techniques & neutrality in genetic algorithms. In *8th International Conference on Evolutionary Computation Theory and Applications ECTA Porto, Portugal* (2106).
- [94] HOFFMEISTER, F., AND BÄCK, T. Genetic self-learning. In *Proceedings of the 1st European Conference on Artificial Life* (1992), The MIT Press, pp. 227–235.
- [95] HOLLAND, J. H. *Adaptation in natural artificial systems*. University of Michigan Press, Ann Arbor, 1975.
- [96] HORDIJK, W. Population flow on fitness landscapes. Master’s thesis, Erasmus University, Rotterdam, 1994.
- [97] HORDIJK, W. A measure of landscapes. Tech. Rep. Technical report SFI Preprint 95-05-049, Santa Fe Institute., 1995.
- [98] HORN, J., HORN, J., AND GOLDBERG, D. E. Genetic algorithm difficulty and the modality of fitness landscapes. In *Foundations of Genetic Algorithms 3* (1994), Morgan Kaufmann, pp. 243–269.
- [99] HUTTER, M. Fitness uniform selection to preserve genetic diversity. In *In Proc. 2002 Congress on Evolutionary Computation (CEC-2002)* (2002), IEEE, pp. 783–788.

-
- [100] HUYNEN, M. Exploring Phenotype Space Through Neutral Evolution. *Mol Evol* 43 (1996), 165–169.
- [101] HUYNEN, M., STADLER, P. F., AND FONTANA, W. Smoothness within ruggedness: the role of neutrality in adaptation. *Proc Natl Acad Sci U S A.* 93, 1 (1996), 397–401.
- [102] JONES, T. *Evolutionary Algorithms, Fitness Landscapes and Search*. PhD thesis, University of New Mexico, Albuquerque, NM, 1995.
- [103] JONES, T., AND FORREST, S. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In *Proceedings of the Sixth International Conference on Genetic Algorithms* (San Francisco, CA, 1995), L. Eshelman, Ed., Morgan Kaufmann, pp. 184–192.
- [104] K. VASSILEV, V., C. FOGARTY, T., AND MILLER, J. Smoothness, ruggedness and neutrality of fitness landscapes: from theory to application. In *Advances in Evolutionary Computing: Theory and Applications* (01 2003), pp. 3–44.
- [105] KARGUPTA, H., DEB, K. L., AND GOLDBERG, D. E. Ordering genetic algorithms and deception. In *Parallel Problem Solving from Nature PPSN II* (1992), Springer, pp. 47–56.
- [106] KATADA, Y., AND OHKURA, K. Estimating the degree of neutrality in fitness landscapes by the neis standard genetic distance an application to evolutionary robotics . In *Proceedings of the 2006 IEEE Congress on Evolutionary Computation* (2006), IEEE Press, pp. 483–490.
- [107] KATADA, Y., AND OHKURA, K. Analysis on Topologies of Fitness Landscapes with Both Neutrality and Ruggedness Based on Neutral Networks. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation (GECCO '09)* (New York, NY, USA., 2009), ACM, pp. 1855–1856.
- [108] KATADA, Y., OHKURA, K., AND UEDA, K. An approach to evolutionary robotics using a genetic algorithm with a variable mutation rate strategy. In *PPSN* (2004), vol. 3242, LNCS Springer, pp. 952–961.
- [109] KAUFFMAN, S. A. *Origins of Order*. Oxford University Press, Oxford., 1993.
- [110] KENNEDY, J., AND SPEARS, W. Matching algorithms to problems: an experimental test of the particle swarm and some genetic algorithms on the multimodal problem generator. In *Proceedings of the IEEE World Congress on Computational Intelligence*, pp. 74–77, Anchorage: Alaska, May 1998. IEEE. (1998).

-
- [111] KIMURA, M. Evolutionary Rate at the Molecular Level. *Nature* 217, 1 (1968), 624–626.
- [112] KIMURA, M. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, 1983.
- [113] KING, J. L., AND DUKES, T. H. Non-Darwinian Evolution. *Science* 164 (1969), 788–798.
- [114] KNOWLES, J. D., AND WATSON, R. A. On the utility of redundant encodings in mutation-based evolutionary search. In *Parallel Problem Solving from Nature - PPSN VII: 7th International Conference* (Granada, Spain, 2002), Springer Verlag, pp. 88–98.
- [115] KUBALIK, J. Using Genetic Algorithms with Real-coded Binary Representation for Solving Non-stationary Problems. In *Adaptive and Natural Computing Algorithms*, B. Ribeiro, R. F. Albrecht, A. Dobnikar, D. W. Pearson, and N. Steele, Eds. Springer Vienna, 2005, pp. 222–225.
- [116] LEFORT, V., KNIBBE, C., BESLON, G., AND FAVREL, J. A bio-inspired genetic algorithm with a self-organizing genome: the rbf-gene model. In *Genetic and Evolutionary Computation Gecco* (2004).
- [117] LEHRE, P. K., AND HADDOW, P. C. Accessibility between neutral networks in indirect genotype-phenotype mappings. In *IEEE Congress on Evolutionary Computation (CEC 2005)* (Edinburgh, UK, 2005), IEEE Press, pp. 419–426.
- [118] LEHRE, P. K., AND HADDOW, P. C. Phenotypic complexity and local variations in neutral degree. *BioSystems* 87, 2-3 (2006), 233–42.
- [119] LEWONTIN, R. C. *The Genetic Basis of Evolutionary Change*. Columbia University Press., New York and London., 1974.
- [120] LIPSICH, M. Adaptation on rugged landscapes generated by iterated local interactions of neighboring genes. In *Proc. of the Fourth Int. Conf. on Genetic Algorithms* (San Mateo, CA, 1991), R. K. Belew and L. B. Booker, Eds., Morgan Kaufmann, pp. 128–135.
- [121] LOBO, J., MILLER, J. H., AND FONTANA, W. Neutrality in technological landscapes. Tech. rep., Santa Fe Working Paper, 2004.
- [122] LOUIS, S. J., AND RAWLINS, G. J. E. Syntactic analysis of convergence in genetic algorithms. In *Foundations of Genetic Algorithms 2* (1992), Morgan Kaufmann, pp. 141–151.

- [123] LOUIS, S. J., AND RAWLINS, G. J. E. Predicting convergence time for genetic algorithms. Tech. Rep. Technical Report TR370, School of Informatics And Computing Indiana University, 1993.
- [124] LOVÁSZ, L. Random walks on graphs: A survey. In *Combinatorics, Paul Erdős is Eighty*, D. Miklós, V. T. Sós, and T. Szőnyi, Eds., vol. 2. János Bolyai Mathematical Society, Budapest, 1993, pp. 1–46.
- [125] MAHFOUD., S. W. Crowding and preselection revisited. In *Parallel problem solving from nature 2* (1992), North-Holland, pp. 27–36.
- [126] MAHFOUD, S. W. *Niching Methods for Genetic Algorithms*. PhD thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 1995. UMI Order No. GAX95-43663.
- [127] MANDERICK, B. The genetic algorithm and the structure of the fitness landscape. *Proc. 4th International Conference on Genetic Algorithms* (1991), 143–150.
- [128] MANNER, R., MAHFOUD, S., AND MAHFOUD, S. W. Crowding and preselection revisited. In *Parallel Problem Solving From Nature* (1992), North-Holland, pp. 27–36.
- [129] MCPHEE, N. F., AND HOPPER, N. J. Analysis of genetic diversity through population history. In *GECCO99: Proceedings of the Genetic and Evolutionary Computation* (1999).
- [130] MIGLINO, O., NOLFI, S., AND PARISI, D. Discontinuity in evolution: How different levels of organization imply pre-adaptation. In *In* (1995), Addison-Wesley.
- [131] MITCHELL, M. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, 1996.
- [132] MITCHELL, M., AND FORREST, S. Genetic algorithms and artificial life. *Artificial Life* 1, 3 (1994), 267–289.
- [133] MITCHELL, M., FORREST, S., AND HOLLAND, J. The royal road for genetic algorithms: Fitness landscapes and ga performance. In *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life* (Cambridge, MA, 1992), MIT Press, pp. 245–254.
- [134] MORRISON, J., AND OPPACHER, F. Maintaining genetic diversity in genetic algorithms through co-evolution. In *Advances in Artificial Intelligence: 12th Biennial Conference of the Canadian Society for Computational Studies of Intelligence, AI'98 Vancouver, BC, Canada, June 18–20, 1998 Proceedings* (Berlin,

- Heidelberg, 1998), R. E. Mercer and E. Neufeld, Eds., Springer Berlin Heidelberg, pp. 128–138.
- [135] MORRISON, R. W., AND DEJONG, K. A. Measurement of Population Diversity. In *In 5th International Conference EA, 2001, volume 2310 of Incs* (2002), Springer.
- [136] MOTOKI, M., AND UEHARA, R. Unique solution instance generation for 3sat. *Research Report, Dept. of Math. and Computing Sciences, Tokyo Inst. of Tech* (1999).
- [137] MOTOKI, T. Calculating the expected loss of diversity of selection schemes. *Evol. Comput.* 10, 4 (Dec. 2002), 397–422.
- [138] MURAO, H., TAMAKI, H., AND KITAMURA, S. A coevolutionary approach to adapt the genotype-phenotype map in genetic algorithms. In *Proceedings of the Evolutionary Computation on 2002. CEC '02. Proceedings of the 2002 Congress - Volume 02* (Washington, DC, USA, 2002), CEC '02, IEEE Computer Society, pp. 1612–1617.
- [139] NEI, M. *Molecular Evolutionary Genetics*. Columbia Univ Press, New York, 1987.
- [140] NEWMAN, M. E. J., AND ENGELHARDT, R. Effects of neutral selection on the evolution of molecular species. Tech. Rep. Working Papers 98-01-001, Santa Fe Institute, 1998.
- [141] NEWMAN, M. E. J., AND ENGELHARDT, R. Effects of selective neutrality on the evolution of molecular species. In *In Proceedings of the Royal Society of London* (1998), vol. 265, pp. 1333–1338.
- [142] NIMWEGEN, E. V., CRUTCHFIELD, J., AND MITCHELL, M. Statistical dynamics of the royal road genetic algorithm. *Theor. Comput. Sci.* 229 (1999), 41–102.
- [143] NIMWEGEN, E. V., CRUTCHFIELD, J., AND HUYNEN, M. Neutral evolution of mutational robustness. *Proc. Natl. Acad. Sci. USA* 96, 17 (1999), 9716–9720.
- [144] NOH, J. D., AND RIEGER, H. Random walks on complex networks. *Physical review letters* 92, 11 (2004), 118701.
- [145] NOLFI, S., MIGLINO, O., AND PARISI, D. Phenotypic plasticity in evolving neural networks. In *From Perception to Action Conference, 1994. Proceedings* (2002), pp. 146–157.

- [146] PELIKAN, M. *Hierarchical Bayesian Optimization Algorithm - Toward a New Generation of Evolutionary Algorithms*, vol. 170 of *Studies in Fuzziness and Soft Computing*. Springer, 2005.
- [147] POLI, R., AND GALVÁN-LÓPEZ, E. The effects of constant and bit-wise neutrality on problem hardness, fitness distance correlation and phenotypic mutation rates. *IEEE Transactions on Evolutionary Computation* 16, 2 (April 2012), 279–300.
- [148] RAVEN, M. J., AND BEDAU, M. A. General framework for evolutionary activity. In *ECAL* (2003), pp. 676–685.
- [149] REIDYS, C. M., AND STADLER, P. F. Neutrality in fitness landscapes. *Applied Mathematics and Computation* 117 (1998), 187–207.
- [150] REIDYS, C. M., STADLER, P. F., AND SCHUSTER, P. Generic properties of combinatorial maps - neutral networks of rna secondary structures. *Bull. Math. Biol.* 59 (1995), 339–397.
- [151] RIDLEY, M. *Evolution*. Blackwell Publishing, Hoboken, USA, 2003.
- [152] RONALD, S., ASENSTORFER, J., AND VINCENT, M. Representational redundancy in evolutionary algorithms. In *Evolutionary Computation, 1995., IEEE International Conference on* (Nov 1995), vol. 2, pp. 631–636 vol.2.
- [153] ROTHLAUF, F. *Representations for Genetic and Evolutionary Algorithms*. Springer-Verlag, 2002.
- [154] ROTHLAUF, F. *Representations for genetic and evolutionary algorithms*, second ed. Springer-Verlag, pub-SV:adr, 2006. First published 2002, 2nd edition available electronically.
- [155] ROTHLAUF, F. Representations for evolutionary algorithms. In *GECCO (Companion)* (2012), pp. 2535–2556.
- [156] ROTHLAUF, F., AND GOLDBERG, D. E. Redundant representations in evolutionary computation. *Evolutionary Computation* 11, 4 (2003), 381–415.
- [157] SALOMON, R. Reevaluating genetic algorithm performance under coordinate rotation of benchmark functions. *BioSystems* 39(3) (1996), 263–278.
- [158] SCHAFFER, J. D., AND ESHELMAN, L. J. On Crossover as an Evolutionary Viable Strategy. In *Proceedings of the 4th International Conference on Genetic Algorithms* (1991), R. Belew and L. Booker, Eds., Morgan Kaufmann, pp. 61–68.

- [159] SCHAFFER, J. D., AND ESHELMAN, L. J. On Crossover as an Evolutionary Viable Strategy. In *Proceedings of the 4th International Conference on Genetic Algorithms* (1991), R. Belew and L. Booker, Eds., Morgan Kaufmann, pp. 61–68.
- [160] SCHAFFER, J. J., CARUANA, R. A., ESHELMAN, L., AND DAS, R. A study of control parameters affecting online performance of genetic algorithms for functional optimization. In *Proc. of the Third Int. Conf. on Genetic Algorithms and their Applications* (San Mateo CA, 1989), J. Schaffer, Ed., Morgan Kaufmann, pp. 51–60.
- [161] SCHEINER, S. M. Genetics and evolution of phenotypic plasticity. In *Annual Review of Ecology and Systematics* (1993), pp. 35–68.
- [162] SCHOENAUER, M., AND MICHALEWICZ, Z. Evolutionary computation, 1997.
- [163] SCHUSTER, P. Genotypes with Phenotypes: Adventures in an RNA toy world. *Biophys Chem* 66, 2 (1997), 75–110.
- [164] SCHUSTER, P., FONTANA, W., STADLER, P. F., AND HOFACKER, I. L. From sequences to shapes and back: a case study in RNA secondary structures. *Proc R Soc Lond B Biol Sci* 255, 1344 (Mar. 1994), 279–284.
- [165] SHACKLETON, M. A., SHIPMAN, R., AND EBNER, M. An investigation of redundant genotype-phenotype mappings and their role in evolutionary search. In *Proceedings of the International Congress on Evolutionary Computation (CEC 2000)* (2000), IEEE Press, pp. 493–500.
- [166] SHANG, Y.-W., AND QIU, Y.-H. A note on the extended rosenbrock function. *Evolutionary Computation* 14, 1 (2006), 119–126.
- [167] SHIPMAN, R. Genetic Redundancy: Desirable or Problematic for Evolutionary Adaption. In *Proceedings of the 4th international Conference on Artificial Neural Networks and Genetic Algorithms (ICANNGA '99)* (Berlin, 1999), A. Dobnikar, N. Steele, D. W. Pearson, and R. F. Albrecht, Eds., Springer, pp. 337–344.
- [168] SHIPMAN, R., SHACKLETON, M., EBNER, M., AND WATSON, R. Neutral search spaces for artificial evolution: A lesson from life. In *Proceedings of the Seventh International Conference on Artificial Life* (2000), MIT Press, pp. 162–169.
- [169] SHIPMAN, R., SHACKLETON, M., AND HARVEY, I. The use of neutral genotype-phenotype mappings for improved evolutionary search. *BT Technology Journal* 18 (October 2000), 103–111.

- [170] SMITH, R. E., AND SMITH, J. E. New methods for tunable, random landscapes. In *Foundations of Genetic Algorithms 6* (2001), Morgan Kaufmann.
- [171] SMITH, T., HUSBANDS, P., LAYZELL, P., AND O'SHEA, M. Fitness landscapes and evolvability. *Evol. Comput.* 10 (March 2002), 1–34.
- [172] SMITH, T., HUSBANDS, P., AND O'SHEA, M. Neutral networks and evolvability with complex genotype-phenotype mapping. In *Proceedings of the 6th European Conference on Advances in Artificial Life, ECAL 01* (London, UK, 2001), Springer-Verlag, pp. 272–281.
- [173] SMITH, T., HUSBANDS, P., AND O'SHEA, M. Neutral networks in an evolutionary robotics search space. In *Proceedings on Evolutionary Computation 2001* (2001), IEEE Press, pp. 136–145.
- [174] SPEARS, W. Crossover or mutation. In *Spears, W.M. (1992) Crossover or Mutation?. In Whitley (ed.) Foundations of Genetic Algorithms-2. 221237.* (1992).
- [175] SPEARS, W. M., DE JONG, K. A., BÄCK, T., FOGEL, D. B., AND DE GARIS, H. An overview of evolutionary computation. In *Proceedings of the European Conference on Machine Learning (ECML-93)* (Vienna, Austria, 1993), P. B. Brazdil, Ed., vol. 667, Springer Verlag, pp. 442–459.
- [176] STADLER, B. M. R., STADLER, P. F., WAGNER, G., AND FONTANA, W. The topology of the possible: Formal spaces underlying patterns of evolutionary change. *Journal of Theoretical Biology* 2, 213 (2001), 241–274.
- [177] STADLER, P. F. Towards a theory of landscapes. In *Complex Systems and Binary Networks: Guanajuato Lectures Held at Guanajuato, México 16–22 January 1995* (Berlin, Heidelberg, 1995), R. López-Peña, H. Waelbroeck, R. Capovilla, R. García-Pelayo, and F. Zertuche, Eds., Springer Berlin Heidelberg, pp. 78–163.
- [178] STADLER, P. F. Genotype-phenotype maps. *Biological Theory* 2 (2006).
- [179] STEPHENS, C. R., AND WAELBROECK, H. Effective degrees of freedom in genetic algorithms and the block hypothesis. In *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)* (San Francisco, CA, 1997), T. Bäck, Ed., Morgan Kaufmann.
- [180] STEPHENS, C. R., AND WAELBROECK, H. Schemata evolution and building blocks. *Evolutionary Computation* 7, 2 (1999), 109–124.
- [181] TATE, D., DAVID, M., AND SMITH, A. Expected allele coverage and the role of mutation in genetic algorithms, 1993.

- [182] TOUSSAINT, M. Demonstrating the evolution of complex genetic representations: An evolution of artificial plants. In *In Proceedings of the 2003 Genetic and Evolutionary Computation Conference (GECCO 2003)* (2003), Springer-Verlag, pp. 86–97.
- [183] TOUSSAINT, M. *The Evolution Of Genetic Representations And Modular Adaptation*. PhD thesis, Institut für Neuroinformatik, Ruhr-Universität Bochum, Germany, 2003.
- [184] TOUSSAINT, M. On the evolution of phenotypic exploration distributions. In *Foundations of Genetic Algorithms 7 (FOGA VII)* (2003), K. A. DeJong, R. Poli, and J. Rowe, Eds., Morgan Kaufmann, pp. 169–182.
- [185] TOUSSAINT, M., AND IGEL, C. Neutrality: A Necessity for Self-adaptation. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2002)* (2002), pp. 1354–1359.
- [186] VASSILEV, V. K., FOGARTY, T. C., AND MILLER, J. F. Information characteristics and the structure of landscapes. *Evol. Comput.* 8 (2000), 31–60.
- [187] VAVAK, F., FOGARTY, T. C., AND JUKES, K. A genetic algorithm with variable range of local search for tracking changing environments. *Proc. 4th Conf. Parallel Problem Solving from Nature, (Lecture Notes in Computer Science) 1141* (1996), 376 – 385.
- [188] VAVAK, F., JUKES, K., AND FOGARTY, T. C. Learning the local search range for genetic optimization in nonstationary environments. In *Proc. 4th IEEE Conf. Evolutionary Computation* (Piscataway, NJ, 1997), IEEE Press, pp. 355 – 360.
- [189] VÉREL, S., COLLARD, P., AND CLERGUE, M. Where are bottlenecks in NK fitness landscapes? In *IEEE International Congress on Evolutionary Computation CEC 2003* (Piscataway, NJ, 2003), IEEE Press, pp. 273–280.
- [190] VOSE, M. D. *The Simple Genetic Algorithm: Foundations and Theory*. MIT Press, Cambridge, MA, USA, 1998.
- [191] WAGNER, A. Robustness, evolvability and neutrality. *FEBS Letters* 579, 8 (2005), 1772–1778.
- [192] WAGNER, G. P., AND ALTENBERG, L. Complex adaptations and the evolution of evolvability. *Evolution* 50, 3 (1996), 967–976.
- [193] WAND, A. Dynamic activation of protein function: a view emerging from nmr spectroscopy. *Nature Structural & Molecular Biology* 8, 11 (2001), 926–931.

- [194] WATSON, R. A., HORNBY, G., AND POLLACK, J. B. Modeling building-block interdependency. In *Proceedings of the 5th International Conference on Parallel Problem Solving from Nature* (London, UK, UK, 1998), PPSN V, Springer-Verlag, pp. 97–108.
- [195] WEICKER, K., AND WEICKER, N. Burden and benefits of redundancy. In *Foundations of Genetic Algorithms 6* (San Francisco, 2000), Morgan Kaufmann, pp. 313–333.
- [196] WHITLEY, D. The genitor algorithm and selection pressure: Why rank-based allocation of reproductive trials is best. In *Proceedings of the Third International Conference on Genetic Algorithms*, J. D. Schaffer, Ed. Morgan Kaufmann, San Mateo, CA., 1989, pp. 116–121.
- [197] WHITLEY, D., MATHIAS, K., RANA, S., AND DZUBERA, J. Evaluating evolutionary algorithms. *Artificial Intelligence 85* (1996), 245–276.
- [198] WHITLEY, D. L. Fundamental Principles of Deception in Genetic Search. In *Foundations of Genetic Algorithms*, G. J. Rawlins, Ed. Morgan Kaufmann, San Mateo, CA, 1991, pp. 221–241.
- [199] WHITLEY, D. L. A genetic algorithm tutorial. *Statistics and Computing 4* (1994), 65–85.
- [200] WHITLEY, D. L., STARKWEATHER, T., AND FUQUAY, D. Scheduling problems and traveling salesmen: The genetic edge recombination operator. In *Proceedings of the 3rd International Conference on Genetic Algorithms* (San Francisco, CA, USA, 1989), Morgan Kaufmann Publishers Inc., pp. 133–140.
- [201] WILKE, C. O., WANG, J. L., OFRIA, C., LENSKI, R. E., AND ADAMI, C. Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature 412* (2001), 331–3.
- [202] WOLPERT, D. H., AND MACREADY, W. G. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation 1*, 1 (April 1997), 67–82.
- [203] WRIGHT, S. The Role of Mutation, Inbreeding, Crossbreeding and Selection in Evolution. In *Proceedings of the Sixth International Congress on Genetics* (1932), D. Jones, Ed., vol. 1, pp. 356–366.
- [204] WU, A. S., AND GARIBAY, I. The proportional genetic algorithm: Gene expression in a genetic algorithm. In *Genetic Programming and Evolvable Hardware* (2002).

-
- [205] YU, T., AND MILLER, J. Neutrality and the evolvability of boolean function landscape. In *Genetic Programming, Proceedings of EuroGP2001, volume 2038 of LNCS* (2001), Springer-Verlag, pp. 204–217.
- [206] YU, T., AND MILLER, J. Finding needles in haystacks is not hard with neutrality. In *Proceedings of the Fifth European Conference on Genetic Programming (EuroGP-2002), volume 2278 of LNCS* (2002), Springer-Verlag, pp. 13–25.
- [207] YU, T., AND MILLER, J. The role of neutral and adaptive mutation in an evolutionary search on the onemax problem. *Blackwell Science Ltd, Global Change Biology* 6 (2002), 291–300.