

# **Rules of Engagement: Competitive Coevolutionary Dynamics in Computational Systems**

**by**

*John Peter Carlidge*

**Submitted in accordance with the requirements  
for the degree of Doctor of Philosophy.**



**The University of Leeds  
School of Computing**

**June 2004**

**The candidate confirms that the work submitted is his own and that the appropriate credit has been given where reference has been made to the work of others. This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.**

# Abstract

Given that evolutionary biologists have considered coevolutionary interactions since the dawn of Darwinism, it is perhaps surprising that coevolution was largely overlooked during the formative years of evolutionary computing. It was not until the early 1990s that Hillis' seminal work thrust coevolution into the spotlight.

Upon attempting to evolve fixed-length sorting networks, a problem with a long and competitive history, Hillis found that his standard evolutionary algorithm was producing sub-standard networks. In response, he decided to reciprocally evolve a population of test-lists against the sorting network population; thus producing a coevolutionary system. The result was impressive; coevolution not only outperformed evolution, but the best network it discovered was only one comparison longer than the best-known solution. For the first time, a coevolutionary algorithm had been successfully applied to problem-solving.

Pre-Hillis, the shortcomings of standard evolutionary algorithms had been understood for some time: whilst defining an adequate fitness function can be as challenging as the problem one is hoping to solve, once achieved, the accumulation of fitness-improving mutations can push a population towards local optima that are difficult to escape. Coevolution offers a solution. By allowing the fitness of each evolving individual to vary (through competition) with other reciprocally evolving individuals, coevolution removes the requirement of a fitness *yardstick*. In conjunction, the reciprocal adaptations of each individual begin to erode local optima as soon as they appear.

However, coevolution is no panacea. As a problem-solving tool, coevolutionary algorithms suffer from some debilitating dynamics, each a result of the relative fitness assessment of individuals. In a single-, or multi-, population competitive system, coevolution may stabilize at a suboptimal equilibrium, or mediocre stable state; analogous to the traditional problem of local optima. Populations may become highly specialized in an unanticipated (and undesirable) manner; potentially resulting in brittle solutions that are fragile to perturbation. The system may cycle; producing dynamics similar to the children's game rock-paper-scissors. Disengagement may occur, whereby one population out-performs another to the extent that individuals cannot be discriminated on the basis of fitness alone; thus removing selection pressure and allowing populations to drift. Finally, coevolution's relative fitness assessment renders traditional visualization techniques (such as the graph of fitness over time) obsolete; thus exacerbating each of the above problems.

This thesis attempts to better understand and address the problems of coevolution through the design and analysis of simple coevolutionary models. "Reduced virulence"—a novel technique specifically designed to tackle disengagement—is developed. Empirical results demonstrate the ability of reduced virulence to combat disengagement both in sim-

ple and complex domains, whilst outperforming the only known competitors. Combining reduced virulence with diversity maintenance techniques is also shown to counteract mediocre stability and over-specialization.

A critique of the CIAO plot—a visualization technique developed to detect coevolutionary cycling—highlights previously undocumented ambiguities; experimental evidence demonstrates the need for complementary visualizations. Extending the scope of visualization, a first exploration into coevolutionary steering is performed; a technique allowing the user to interact with a coevolutionary system during run-time. Using a simple model incorporating reduced virulence, the coevolutionary steering demonstration highlights the future potential of such tools for both research and education.

The role of neutrality in coevolution is discussed in detail. Whilst much emphasis is placed upon neutral networks in the evolutionary computation literature, the nature of coevolutionary neutrality is generally overlooked. Preliminary ideas for modelling coevolutionary neutrality are presented.

Finally, whilst this thesis is primarily aimed at a computing audience, strong reference to evolutionary biology is made throughout. Exemplifying potential crossover, the CIAO plot, a tool previously unused in biology, is applied to a simulation of *E. Coli*, with results confirming empirical observations of real bacteria.

# Acknowledgements

Thankfully, throughout my PhD I have received tremendous help and support from many colleagues, friends and family: to each and every one, I am eternally grateful.

For love, support and patience throughout, I want to thank Helen Pickup and Mum. For technical help, I want to give particular thanks to Martin Thompson and Lee Bulmer: without Martin's graphical expertise and computer support, the steering interface would have failed to leave the drawing board; without Lee's coding skills, I'd still be sorting through reams of data by hand. I would like to thank Steven Walker for preliminary CIAO plot investigations (Walker, 2000) and Dave Gordon for drawing my attention to similarities between reduced virulence and the Laffer Curve. Special thanks goes to Dave Harris, Simon Lessels and Hywel Williams for guidance, support and general input throughout; nobody else has had to endure my ranting quite so much! Thanks to the Biosystems group at the University of Leeds, both past and present, for being a great group to work with; and the Biosystems reading group for many hours of stimulating conversation and drinking. Thanks to all those from outside Leeds whose conversations and opinions have helped to direct my research, including (especially) Richard Watson, Jordan Pollack, Dave Cliff, Peter Todd and Geoffrey Miller. A big thanks to Neil Currums for giving me time whenever I needed it, Tom Bysouth for accommodation during my many last-minute visits and Tom Barnes-Lawrence for access to his computer. Finally, none deserves acknowledgement more than Seth Bullock, my supervisor and mentor, without whom this thesis would not have been possible: thanks, Seth.

To the one who would have enjoyed it most.

# Declarations

Some parts of the work presented in this thesis have been published in the following articles:

**Cartlidge, J. & Bullock, S. (2002)**, “Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization”. In *Congress on Evolutionary Computation*, pp. 1420–1425. IEEE Press, Piscataway NJ.

**Bullock, S., Cartlidge, J. & Thompson, M. (2002)**, “Prospects for computational steering of evolutionary computation”, in Bilotta, E., et al. (Eds.) *Workshop Proceedings, Artificial Life VIII*, pp 131–137. MIT Press, Cambridge MA.

**Cartlidge, J. & Bullock, S. (2003)**, “Caring Versus Sharing: How to Maintain Engagement and Diversity in Coevolving Populations”, in Banzhaf, W., Christaller, T., Dittrich, P., Kim, J., & Ziegler, J. (Eds.) *Seventh European Conference on Artificial Life*, pp 299–308. Springer-Verlag, Berlin, Germany.

**Cartlidge, J. & Bullock, S. (2004a)**, “Combating coevolutionary disengagement by reducing parasite virulence”, *Evolutionary Computation*, 12(2), 193–222.

**Cartlidge, J. & Bullock, S. (2004b)**, “Unpicking tartan CIAO plots: Understanding irregular coevolutionary cycling”, *Adaptive Behaviour*, 12(2), 69–92.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Outline . . . . .	1
1.1.1	The Problem . . . . .	2
1.1.2	Motivation . . . . .	3
1.1.3	Overview of Results . . . . .	3
1.1.4	Methodology . . . . .	4
1.2	How to Read this Thesis . . . . .	5
1.3	A History of the Coevolutionary Perspective . . . . .	5
1.3.1	Interspecific Interaction . . . . .	6
1.3.2	The Pre-Darwinian Perspective . . . . .	7
1.3.3	Darwin’s Entangled Bank . . . . .	8
1.3.4	Challenges . . . . .	10
1.3.5	Post-Synthesis Developments . . . . .	11
1.3.5.1	The Evolutionary Synthesis . . . . .	11
1.3.5.2	New Techniques: A New Perspective . . . . .	13
1.3.5.3	The Coining of Coevolution . . . . .	16
1.4	Summary . . . . .	18
<b>2</b>	<b>Review of Evolutionary Computation</b>	<b>19</b>
2.1	Artificial Evolution . . . . .	19
2.2	Genetic Algorithms (GAs) . . . . .	21
2.2.1	A Simple Generational GA . . . . .	22
2.2.2	Worked Example . . . . .	23
2.2.3	Modes of Variation . . . . .	24
2.2.3.1	Mutation . . . . .	24
2.2.3.2	Recombination . . . . .	25
2.2.4	Assessment and Selection . . . . .	26
2.3	Realism . . . . .	29

2.3.1	Fixed Population Size . . . . .	31
2.3.2	Genotype-Phenotype Mapping . . . . .	31
2.3.3	Static Environment . . . . .	32
2.4	Summary . . . . .	33
<b>3</b>	<b>Coevolutionary Computation</b>	<b>34</b>
3.1	Background . . . . .	34
3.2	Defining Coevolution . . . . .	36
3.2.1	The Standard Definition of Coevolution . . . . .	37
3.2.2	The Coevolutionary Continuum . . . . .	38
3.2.3	When is it <i>not</i> Coevolution? . . . . .	45
3.2.4	Fitness Propensity and Fitness Realization . . . . .	46
3.2.5	A New Definition of Coevolution . . . . .	49
3.2.6	Relativism versus Absolutism . . . . .	49
3.2.7	Moving Mountains: A Dynamic Landscape . . . . .	51
3.2.8	The Nature of Objective Fitness . . . . .	53
3.3	The Success Stories . . . . .	54
3.4	Critique . . . . .	56
3.4.1	Visualization . . . . .	56
3.4.2	Overspecialisation . . . . .	59
3.4.3	Cycling . . . . .	61
3.4.4	Mediocre Stability . . . . .	67
3.4.5	Disengagement . . . . .	68
3.4.5.1	Classifying Disengagement . . . . .	71
3.4.5.2	Coevolutionary Neutrality . . . . .	73
3.5	Classroom Analogies . . . . .	74
3.6	Summary . . . . .	76
<b>4</b>	<b>Research Statement</b>	<b>78</b>
4.1	Objective . . . . .	78
4.2	Justification . . . . .	79
4.3	Prospects . . . . .	80
<b>5</b>	<b>Reduced Parasite Virulence</b>	<b>81</b>
5.1	Rationale . . . . .	81
5.1.1	Here Today, Gone Tomorrow . . . . .	81
5.1.2	Don't Bite the Hand that Feeds You . . . . .	82

5.2	Analogies . . . . .	83
5.2.1	Parasite Virulence . . . . .	83
5.2.1.1	Myxomatosis: An Illustrative Example . . . . .	85
5.2.1.2	A Case of Group Selectionism? . . . . .	86
5.2.2	The Laffer Curve . . . . .	87
5.3	Implementation . . . . .	89
5.3.1	Maintaining Relative Fitness Diversity . . . . .	90
5.3.2	Asymmetry . . . . .	91
5.4	Summary . . . . .	91
<b>6</b>	<b>Combating Disengagement</b>	<b>92</b>
6.1	The Numbers Game . . . . .	93
6.1.1	Background . . . . .	93
6.1.2	Investigating Engagement . . . . .	94
6.1.2.1	Set Up . . . . .	94
6.1.2.2	Results . . . . .	95
6.1.3	Further Investigations . . . . .	97
6.1.3.1	Comparing the Phantom Parasite . . . . .	98
6.1.3.2	The Engagement-Optimality Trade Off . . . . .	100
6.1.3.3	Re-engagement . . . . .	101
6.1.3.4	Random Parasites: Canonical Evolution . . . . .	101
6.1.4	Conclusions . . . . .	103
6.2	The Coin Toss Game . . . . .	103
6.2.1	Diversity Maintenance . . . . .	103
6.2.2	Set Up . . . . .	104
6.2.3	Results . . . . .	107
6.2.4	Caring versus Sharing . . . . .	110
6.3	Sorting Networks . . . . .	111
6.3.1	Background . . . . .	111
6.3.2	Set Up . . . . .	112
6.3.3	Results . . . . .	113
6.3.3.1	Non-Spatial Model . . . . .	113
6.3.3.2	Spatial Model . . . . .	114
6.4	Dynamic Parasite Virulence . . . . .	116
6.4.1	Rationale . . . . .	116
6.4.2	Measuring Engagedness . . . . .	116



6.4.3	Implementation . . . . .	117
6.4.4	Investigations . . . . .	118
6.4.4.1	The Numbers Game . . . . .	118
6.4.4.2	Sorting Networks . . . . .	120
6.5	Summary . . . . .	122
<b>7</b>	<b>Visualizing Coevolutionary Dynamics</b>	<b>123</b>
7.1	Background . . . . .	124
7.1.1	EC Visualization . . . . .	124
7.1.2	Visualizing Coevolutionary Systems . . . . .	125
7.2	CIAO Plots . . . . .	126
7.2.1	Plotting CIAO Contests . . . . .	128
7.2.2	Progress and Cycling . . . . .	128
7.2.3	Unpicking Tartan CIAO Plots . . . . .	129
7.3	Rock–Paper–Scissors: An Intransitive Domain . . . . .	131
7.3.1	Study 1: A Simple Encoding . . . . .	131
7.3.1.1	Feature Detection . . . . .	132
7.3.1.2	Results . . . . .	133
7.3.2	Study 2: A Complex Encoding . . . . .	136
7.3.2.1	Results . . . . .	138
7.3.3	Study 3: Random Drift . . . . .	142
7.3.4	Understanding Irregular Coevolutionary Cycling . . . . .	145
7.4	Study 4: A Simulation of <i>E. coli</i> . . . . .	148
7.4.1	Replication . . . . .	148
7.4.2	Results . . . . .	149
7.5	Revisiting the Coin Toss Game . . . . .	153
7.6	Steering Coevolution Through Interactive Visualization . . . . .	154
7.6.1	Prospects for Computational Steering . . . . .	156
7.6.2	Proof of Concept . . . . .	157
7.6.2.1	Steering the Numbers Game . . . . .	157
7.6.2.2	Example Output . . . . .	159
7.6.2.3	Future Work . . . . .	162
7.7	Summary . . . . .	165
<b>8</b>	<b>Transient Neutrality</b>	<b>167</b>
8.1	Modelling Neutrality With <i>NK</i> Landscapes . . . . .	169
8.2	Towards an <i>NKCqd</i> Landscape . . . . .	172

8.3	Summary . . . . .	174
<b>9</b>	<b>Conclusions</b>	<b>175</b>
9.1	Summary of Contributions . . . . .	176
9.2	Limitations of the Thesis . . . . .	176
9.3	Future Research . . . . .	177
9.4	Verification . . . . .	178
	<b>Bibliography</b>	<b>179</b>

# List of Figures

2.1	Example GA Output . . . . .	24
3.1	The Coevolutionary Continuum . . . . .	39
3.2	Fitness Propensity and Fitness Realization . . . . .	48
3.3	Invasion and Extinction Cycles in <i>Anolis</i> Lizards . . . . .	63
3.4	The Phantom Parasite and $\Phi$ Functions . . . . .	70
5.1	The Laffer Curve . . . . .	88
5.2	Reduced Parasite Virulence . . . . .	89
6.1	Numbers Game Results: Bias 0.75 . . . . .	95
6.2	Numbers Game Results: Bias 0.9 . . . . .	96
6.3	The Effects of Mutation on Engagement . . . . .	99
6.4	The Engagement of Random Parasites . . . . .	102
6.5	The Coin Toss Game . . . . .	106
6.6	Coin Toss Results . . . . .	108
6.7	A Four-Element Sorting Network . . . . .	112
6.8	Sorting Network Results: Non-Spatial Asexual . . . . .	114
6.9	Sorting Network Results: Spatial Asexual . . . . .	115
6.10	Dynamic Virulence in the Numbers Game . . . . .	119
6.11	Dynamic Virulence and Engagement in the Numbers Game . . . . .	120
6.12	Sorting Network Results: Dynamic Virulence . . . . .	121
7.1	CIAO Plot Schematic . . . . .	128
7.2	Idealised CIAO Plots . . . . .	129
7.3	“Tartan” CIAO Plot Schematic . . . . .	130
7.4	RPS Study 1 Results: CIAO Plots . . . . .	133
7.5	RPS Study 1 Results: Event Plot . . . . .	133
7.6	RPS Study 1 Results: Probability Graphs . . . . .	134
7.7	RPS Study 1 Results: State Transition Table . . . . .	135

7.8	RPS Study 2: Example FSM Encoding . . . . .	136
7.9	RPS Study 2 Results: CIAO Plots . . . . .	138
7.10	RPS Study 2 Results: Event Plot . . . . .	138
7.11	RPS Study 2 Results: Probability Graph . . . . .	139
7.12	RPS Study 2 Results: State Transition Table . . . . .	141
7.13	RPS Study 2: Enlarged CIAO Plot . . . . .	142
7.14	RPS Study 3 Results: CIAO Plots . . . . .	143
7.15	RPS Study 3 Results: Event Plot . . . . .	143
7.16	RPS Study 3 Results: Probability Graphs . . . . .	144
7.17	RPS Study 3 Results: State Transition Table . . . . .	145
7.18	RPS Study: Fitness Variation . . . . .	147
7.19	<i>E. coli</i> Simulation: Grid Output . . . . .	149
7.20	<i>E. coli</i> Simulation: CIAO Plots . . . . .	150
7.21	<i>E. coli</i> Simulation: Event Plot . . . . .	151
7.22	<i>E. coli</i> Simulation: Probability Graphs . . . . .	151
7.23	<i>E. coli</i> Simulation: State Transition Table . . . . .	152
7.24	Coin Toss Results: CIAO plots . . . . .	155
7.25	Steering Interface . . . . .	158
7.26	Example Fitness Distributions . . . . .	159
7.27	Steering Interface: Screen Shot 1 . . . . .	160
7.28	Steering Interface: Screen Shot 2 . . . . .	160
7.29	Steering Interface: Screen Shot 3 . . . . .	161
7.30	Steering Interface: Screen Shot 4 . . . . .	161
7.31	Steering Interface: Screen Shot 5 . . . . .	163
7.32	Steering Interface: Screen Shot 6 . . . . .	163
7.33	Steering Interface: Screen Shot 7 . . . . .	164
7.34	Steering Interface: Screen Shot 8 . . . . .	164
8.1	Example $NK$ , $NKp$ and $NKq$ Landscapes . . . . .	171
8.2	A Partial $NKCqd$ Landscape . . . . .	173

# Chapter 1

## Introduction

---

### 1.1 Outline

Simultaneously decimating the creationist paradigm and the common misconception of humans as distinct from (and somehow *better* than) the rest of the animal kingdom, the publication of Charles Darwin's (1859) *On the origin of species* induced one of the greatest intellectual revolutions in the history of mankind. Today, evolutionary theory pervades virtually every facet of modern biology. However, some aspects of evolution, such as the evolution of interactions between species (or coevolution), are much less pervasive than more "traditional" evolutionary topics, such as the study of genetics.

Following the publication of *Origin*, coevolution was largely ignored by evolutionary biologists for more than 100 years. However, in recent decades, interest in (and understanding of) coevolutionary systems has grown significantly. Since the evolutionary synthesis<sup>1</sup> (1936-1947), the introduction of novel tools, such as evolutionary game theory and the digital computer, have made a major contribution to the development of coevolutionary theory. Yet, whilst much progress has been made, coevolutionary biology is undoubtedly still in its infancy.

In the 1960s, the evolutionary process was first harnessed computationally in the form of an evolutionary algorithm (EA). Such algorithms, incorporating populations of simple individuals, can be utilised either as a modelling tool for biologists or as an optimisation

---

<sup>1</sup>Also named the "modern synthesis", this period witnessed a synthesis of the Mendelian and Darwinian paradigms.

tool for complex problem solving. When running an evolutionary algorithm, the population adapts to the specific task they are given; whether that be a computational problem or an abstract model of a natural system. In both domains, evolutionary algorithms have had fair success.

Since 1990, coevolutionary interactions have been added to the standard evolutionary algorithm. Initially introduced as a method of improving the ability of EAs to optimise, coevolutionary algorithms also facilitate the evolutionary modelling of biological interspecies interactions. With potential not only as a powerful problem-solving tool capable of outperforming standard evolutionary algorithms, but also (uniquely) as a tool for sophisticated individual-based modelling of coevolutionary interactions, coevolutionary algorithms are an exciting innovation. Unfortunately, however, the full potential of coevolutionary algorithms, specifically as an engineering tool, has not yet been realised.<sup>2</sup> To address this, several fundamental problems must firstly be confronted.

The primary aim of this thesis is to improve the utility of coevolutionary algorithms for problem solving by furthering current understanding of the dynamics of artificial coevolutionary systems. As a subsidiary goal, it is hoped that the knowledge derived in this thesis will further theoretical understanding of biological coevolution.

### 1.1.1 The Problem

The dynamics of coevolutionary systems are notoriously difficult to understand and analyse. This proves a severe hindrance when attempting to utilise coevolution as a computational tool. As a result, unless comprehension improves, artificial coevolutionary algorithms are unlikely to fulfil their potential as a “general purpose” optimisation tool.

This thesis aims to begin to address these issues by improving the understanding (and control) of the coevolutionary dynamics of computational systems, particularly two-population antagonistic (or competitive) systems. Where relevant, related insights into biological coevolution will also be considered.

Specifically, this thesis has four aims (unequal emphasis; see chapter 4 for a more detailed discussion):

1. To clarify definitions of coevolution in biological and artificial systems.
2. (a) To further understanding of *disengagement*, a deleterious phenomenon that hinders coevolutionary progress.

---

<sup>2</sup>Where the term “engineering” is used in this thesis, reference is made to the engineering side of artificial life (the autonomous robotics community, for example) and not engineering in general (electronic, chemical, civil, etc.). This thesis is focused upon developing and understanding coevolutionary algorithms so that they may be more productively used in the future; it does not attempt to solve problems picked from the field.

- (b) To produce a novel algorithm with which to tackle disengagement.
3. (a) To further understanding of the specific challenges facing appropriate visualization of coevolution.  
(b) To critique present visualization techniques and suggest novel alternatives.
4. To further understanding of neutrality in coevolutionary systems, and the relationship between neutrality and other coevolutionary phenomena.

### 1.1.2 Motivation

This thesis is primarily motivated by the desire to improve the ability of coevolutionary algorithms to perform as a problem solving technique. Having achieved success in a variety of applications over the past few decades, evolutionary computation nevertheless suffers from some practical problems that are difficult (if not impossible) to overcome; two of the most serious include premature (sub-optimal) convergence of solutions, and difficulties in defining an adequate assessment of fitness. Coevolutionary computation, able to circumvent some of the problems associated with standard evolutionary approaches, has been proposed as an alternative. However, the introduction of coevolution introduces several new problems each associated with relatively poorly understood coevolutionary dynamics.

If successful, this thesis should further understanding of coevolutionary dynamics, whilst introducing novel methods to tackle some of coevolution's problems. The future benefits of improving coevolutionary computation are potentially widespread in engineering applications and may also improve biological modelling techniques. In conjunction, many of the questions raised in this thesis have biological analogies. As such, there is a possibility that answering these questions may shed light upon coevolutionary phenomena in the biological world. The possibility of addressing such fundamental scientific questions is a strong motivator.

### 1.1.3 Overview of Results

The major results presented in this thesis are summarised below:

1. Definitions of coevolution, both in biological and artificial systems, are clarified (see chapter 3).
2. The phenomenon of disengagement is classified for the first time (chapter 3). A novel technique designed to combat the effects of disengagement is introduced

- (chapter 5), with empirical evidence suggesting that alternative state-of-the-art approaches are outperformed (chapter 6).
3. A detailed investigation of CIAO plots—a standard coevolutionary visualization—uncovers previously unobserved ambiguity (chapter 7).
  4. A novel *interactive* visualization method incorporating a “steering” mechanism is introduced. Such steering allows run time manipulation of the coevolutionary system (chapter 7).
  5. A discussion of neutrality in coevolutionary systems is undertaken for the first time and a novel framework for investigating coevolutionary neutrality is suggested (chapter 8).

### 1.1.4 Methodology

As part of the general research methodology used throughout this thesis, the evolutionary biology literature is extensively used as a source for ideas and inspiration with which to tackle the problems of artificial coevolution. Having some overlap with coevolutionary computation, it seems plausible to expect that the more mature field of evolutionary biology may contain established knowledge and techniques with potential application to artificial coevolution. A better understanding of the similarities and differences between artificial and biological systems can be gained by integrating knowledge between these fields.

To test and develop preliminary ideas, trivial problem domains are used for early experiments and analysis. Such simplistic models aid transparency, allow tractable results and afford more insight. More complex problem domains are used for further development of work. Whilst simple models allow general dynamics to be observed, more complex problems show whether findings are transferable to domains of “interest” whilst demonstrating the potential application of novel techniques to non-trivial optimisation problems. However, the “big question”, and one that unfortunately remains open, is how to discern the set of interesting problems an algorithm is best suited to.

Finally, once developments in artificial coevolutionary systems have been established, considerations are made as to whether such developments can give insight into biological coevolution. In this way it is hoped that ideas can circulate between evolutionary biology and evolutionary computation.



## 1.2 How to Read this Thesis

Bio-inspired computing, as its name suggests, lies at the interface between computing and biology. Inspired by biological systems, novel computer algorithms, such as the ubiquitous genetic algorithm (see chapter 2), are developed as a tool for problem solving, whilst developments in computational understanding filter back into biology via innovative modelling techniques, such as agent-based simulations. As such, it is possible to approach bio-inspired computing from two perspectives; either that of a biologist, or that of a computer scientist.

However, it is unwise to separate biology and computing, or biological modelling and problem solving, exclusively. In order to fully appreciate one angle it is important to have a firm grasp of the other. Thus, although this thesis is primarily concerned with the “computational” perspective, attention will also be focused, where necessary, from the perspective of biology.

To aid comprehension, it is important for the reader to understand which perspective is taken at any one time; “optimisation” for instance, has a very different connotation when taken from a computational rather than a biological viewpoint. Thus, when a biological perspective is taken, it shall be, where appropriate, highlighted by framing the text within a box, as such:

This text is highlighted by a framing box in order to indicate that the reader should be taking a biological perspective. This notation will only be used where appropriate.

Due to the strong coupling between biology and computing, it is intended that, before introducing coevolutionary computation, the reader shall receive a historical account of coevolution in evolutionary biology. Integral to understanding the development of artificial coevolution whilst anticipating possible future developments, the following section outlines a brief history of evolutionary biology from a coevolutionary perspective.

## 1.3 A History of the Coevolutionary Perspective

Within evolutionary biology the study of coevolution has historically received relatively little attention; a remarkable fact given how pervasive coevolutionary interactions are in nature. No species has an evolutionary history entirely independent of other biota, yet it was not until 1979 that the first textbook to dedicate an entire chapter to coevolution appeared (Futuyma, 1979). As a result, coevolution is often perceived as ancillary to the “core” evolutionary framework.

In this section, I will attempt to address this issue by presenting a brief history of evolutionary biology from a coevolutionary perspective. The reason for this is not to espouse a paradigm shift, or scientific revolution *à la* Kuhn (1970; for an extended discussion see Chalmers, 1999)—it is certainly not my intention to encourage the reader to see *everything* as coevolution—but to highlight that coevolutionary ideas have a long history, traceable to the inception of evolutionary theory. I wish to demonstrate that, far from being anomalous, coevolutionary theory is *integral* to evolutionary biology. The reason why coevolution has been under-represented in evolutionary research is not because it is uninteresting, but rather that the coevolutionary dynamics involved in inter-species interactions are more difficult to analyse than many other facets of evolutionary theory.

### 1.3.1 Interspecific Interaction

Given that all animals eat plants, or other animals, interspecific interactions are an inescapable necessity of the natural world. Species simply *do not* evolve in isolation, but rather, to varying degrees, have an evolutionary history of interactions with other biota. In the majority of situations, the primary driving force behind evolutionary adaptations are interspecific interactions—the biotic environment generally changes at a greater rate than the physical environment. For this reason, much of the diversity and specialisation observable within the natural world is due to interspecific interactions, or coevolution.

Indeed, in the form of endosymbiotic<sup>3</sup> interactions, coevolution is considered directly responsible for one of the most significant transitions in evolutionary history. Margulis (1970) postulates that the emergence of eukaryotes—organisms with nucleated cells—is due to the incorporation of free-living prokaryotes (now represented by mitochondria and chloroplasts) by a host cell. This is no small matter. Without eukaryotes, there would be no plants, animals or fungi.

However, despite the impact that the complex evolved network of interspecific interactions has undoubtedly had upon the evolution of species, there has been a distinct lack of emphasis upon coevolution throughout much of the history of evolutionary biology. The evolution of interspecific interactions has often been perceived as a side-issue.

... generations of students in courses in evolutionary biology were trained without realizing that they were studying only half the problem of the evolution of adaptation and diversity. Or they studied evolving interactions as parts of other courses, giving the impression that the evolution of interactions was a

---

<sup>3</sup>Ehrman (1983) describes endosymbiosis as the “most intimate [of] coevolved systems... in which one of the organisms incorporates the other—an intricate co-habitation of two species in which a symbiont exists in the cells of a host for at least a [proportion] of its partner’s life cycle”.

specialised topic outside the scope of the fundamental questions in evolution (Thompson, 1994, p55).

In an attempt to understand why the evolution of interspecific interactions, or coevolution, was so marginalized for such an extensive period of time, it is necessary to review the history of evolutionary biology from a coevolutionary perspective. However, the most salient feature of such a review is the surprising fact that, whether implicit or explicit, talk of interspecific interactions pervaded the formative years of evolutionary biology.

### 1.3.2 The Pre-Darwinian Perspective

In 1809 the French zoologist Jean-Baptiste Pierre Antoine de Monet, Chevalier de Lamarck introduced the world to his fledgling theory of evolution in *Philosophie Zoologique* (1809). Although Lamarck is now best associated with his discredited theory of heredity—the inheritance of acquired characteristics, or “soft” inheritance—Darwin and other early “evolutionists” acknowledged him as a great zoologist and his ideas as a forerunner of evolution.

Pre-Lamarck, the general consensus was a static world with constant conditions; ignoring the odd miraculous interposition. The greatest achievement of *Philosophie Zoologique* was to consider the world—species and environment—as dynamic. Prior to Darwin, the general consensus was *essentialist*, believing that each species had its own species-specific essence, thus making it theoretically impossible for a species to change or evolve. However:

...since Lamarck believed that species could not become extinct, he concluded that they must undergo constant evolutionary change in order to remain adapted to the changes in their environment (Mayr, 1982, p406).

Unfortunately, Lamarck’s ideas were largely ignored in his own lifetime. In 1829 he died in poverty and obscurity.

Lamarck was the first to think about a dynamic world with species evolving in response to a changing environment, though he never explicitly mentioned the effect that evolving species have upon an environment and the subsequent interspecific interactions (Mayr, 1982, Chapter 8). However, explicit musings upon such interactions predated Lamarck’s evolutionary theory. Whether the environment maintained constancy or not, it was obvious to some that species-interactions (though non-evolutionary) pervaded the natural world, beautifully exhibiting the plan of the Creator. Take the following passage from Joseph Priestly in 1764:

Are not all plants likewise suited to the various kinds of animals which feed upon them?...The various kinds of animals are, again, in a thousand ways adapted to, and formed for, the use of one another (Thompson, 1994, p23).

In 1798 the English economist Thomas Robert Malthus first published his *Essay on Population* (Malthus, 1798), acknowledged six decades later by both Darwin and Wallace as having a profound influence upon their subsequent theories of evolution (Mayr, 1982; Nordenskiöld, 1935). Considering the reasons why populations do not increase geometrically, Malthus concluded that populations are checked by a limited food supply—there is not an infinite abundance of resources. Such scarce resources produce a necessity to out-compete in order to survive: there is a *struggle for existence*. Only the strongest, or the luckiest, survive.

Wherever therefore there is liberty, the power of increase is exerted; and the superabundant effects are repressed afterwards by want of room and nourishment, which is common to animals and plants; and among animals, by becoming the prey of others (Malthus, 1798).

Upon marrying the changing environment and constant species adaptation presented by Lamarck with both the interspecific interactions and mutual adaptations of Priestly (ignoring his teleological viewpoint) and the struggle for existence resulting from competition for scarce resources pronounced by Malthus, one could begin to form a proto-coevolutionary theory. However, it was Charles Lyell, a Geologist and (until 1859) anti-evolutionist, who first stated the importance of interspecific interactions explicitly:

Lyell contributed one important thought, subsequently developed by Darwin: it is not only the physical factors that can cause extinction but also competition from other better-adapted species (Mayr, 1982, p406).

Yet—unexceptional given the era—application was one step ahead of theory. Attempting to avoid the potentially devastating effects of crop disease, early nineteenth century plant breeders began the selective breeding of cereal crops for resistance to disease, particularly wheat rust. Though early breeders would not have considered wheat rust to be an evolving organism, they were nevertheless engaged in a coevolutionary struggle (Barrett, 1983).

### 1.3.3 Darwin's Entangled Bank

The publication of *On the Origin of Species* on November 24, 1859, ignited possibly the greatest scientific revolution of all time, and sounded the death knell of anthropocentrism.

To quote Mayr (2001), “[*Origin*] almost single-handedly effected the secularisation of science”. In the first sentence of the final paragraph of *Origin*, Charles Robert Darwin’s most famous of passages encourages the reader to imagine an entangled bank, evidencing that, from the outset, Darwinian evolution incorporated interspecific interactions and, hence, coevolutionary theories.

It is interesting to contemplate an entangled bank, clothed with many plants of many kinds, with birds singing on the bushes, with various insects flitting about, and with worms crawling through the damp earth, and to reflect that these elaborately constructed forms, so different from each other, and dependent on each other in so complex a manner, have all been produced by laws acting around us (Darwin, 1985, p459).

Darwin explicitly tackles coevolutionary interactions in an earlier chapter of *Origin*—the first real account of the evolutionary process. Given the following sentence, there can be no denying that Darwin understood the influence of interspecific interactions.

Thus I can understand how a flower and a bee might slowly become, either simultaneously or one after the other, modified and adapted in the most perfect manner to each other, by the continued preservation of individuals presenting mutual and slightly favourable deviations of structure (Darwin, 1985, p142).

However, Darwin was not alone. After twenty years of procrastination, Darwin unveiled his theory of evolution in *Origin* after becoming aware that Alfred Russel Wallace had arrived at an almost identical theory. In a letter sent to Darwin in 1858, Wallace presented his own theory of evolution. Amazed by the similarity of Wallace’s theory, and fearing that he would be beaten to publication, Darwin attached some of his own work to Wallace’s letter and forwarded it on to Lyell and Hooker for presentation to the Linnean Society of London (Mayr, 1982). Interestingly, though Wallace had independently arrived at his own evolutionary theory, he was also well aware of the importance of interspecific interactions.

Even the peculiar colours of many animals, especially insects, so closely resembling the soil or the leaves or the trunks on which they habitually reside, are explained on the same principle; for though in the course of ages varieties of many tints may have occurred, *yet those races having colours best adapted to concealment from their enemies would inevitably survive the longest* (Wallace, 1858, original italics).

Immediately following *Origin*, Darwin undertook a coevolutionary study of orchids and moths (see Thompson, 1994). He states:

... there has been a race in gaining length between the nectary of *Angraecum* and the proboscis of certain moths (Darwin, 1862, p202-3).

Yet the first landmark in coevolution came not to Darwin, but to Wallace's friend, the entomologist Henry Walter Bates. Having stayed behind in Brazil after Wallace returned to England, Bates developed his theory of mimicry through the observation of butterflies. Bates observed that each species of the unpalatable heliconid butterflies was associated, where it occurred, with one or more species of edible butterflies which mimicked them in colouration (Mayr, 1982)—Batesian mimicry (Bates, 1862). Bates' work became one of the major initial applications of Darwinian evolution and an outstanding contribution to coevolutionary theory (Thompson, 1994).

In 1879, Fritz Müller contributed the most important extension to the mimicry principle; mutual mimicry may occur between unpalatable, poisonous, or venomous animals such as wasps and snakes—Müllerian mimicry (Müller, 1879). Müllerian mimicry comprised the first quantitative statement of the conditions under which reciprocal evolution may occur (Thompson, 1994).

However, despite the initial success of mimicry theory, interspecific interactions were largely ignored in the final decades of the nineteenth century and the pre-synthesis years of the twentieth. In the following section we attempt to understand why this period saw such a dearth of coevolutionary theory.

### 1.3.4 Challenges

In order to study interspecific interactions in the Victorian era, there were several pragmatic challenges one must overcome. Primarily, empirical evidence was lacking. As Thompson (1994) notes, whereas one could study species' morphologies by having them shipped from far continents, species-interactions were not so easy to attain; they cannot be housed in a museum, they have no comparable morphologies and they rarely leave a direct fossil record. Thus, in order to gain a significant understanding of a particular interspecific interaction, one must directly observe the coevolutionary system *in situ*. Unfortunately, this may involve extensive periods in inhospitable conditions, such as those faced by Bates in the Brazilian rainforests during his eleven years of study between 1848-1859. Little wonder then, that species interactions were often overlooked.

Secondly, interspecific interactions simply include an extra level of complexity over and above that found within a single species. Explaining morphologies on the basis of

physical conditions is a much easier task than attempting to understand complex, evolving interspecific interactions, that can often be counter-intuitive. Given that nineteenth century Darwinians faced the challenge of trying to convince and convert anti-evolutionist essentialists, one can imagine why explanations of particular animal and plant morphologies would appear more appropriate than descriptions of complex networks of interspecific interactions. Much of Victorian society was not quite ready for man's relation to primates, let alone coevolutionary theory.

Finally, even if empirical evidence is collected and a receptive audience is discovered, one major challenge still persists—how to interpret such evidence. Complex interspecific interactions are very hard to analyse and interpret. In the pre-synthesis years, few tools existed, making it virtually impossible to study many coevolutionary systems adequately. Coevolution was often not overlooked because of lack of interest, but because it was beyond the scope of contemporary science. Fortunately, the greatest biological revolution of the twentieth century—the evolutionary synthesis—would change this scenario.

### 1.3.5 Post-Synthesis Developments

The term “evolutionary synthesis” was introduced by [Huxley (1942)] to designate the general acceptance of two conclusions: gradual evolution can be explained in terms of small genetic changes (“mutations”) and recombination, and the ordering of this genetic variation by natural selection; and the observed evolutionary phenomena, particularly macroevolutionary processes and speciation, can be explained in a manner that is consistent with the known genetic mechanisms (Mayr & Provine, 1980).

#### 1.3.5.1 The Evolutionary Synthesis

During the initial decades of the twentieth century, two seemingly incommensurable paradigms (as defined by Kuhn, 1970) dominated evolutionary biology; the naturalist camp and the geneticist camp.<sup>4</sup> Despite all that had been learned since *Origin*, the camps disagreed almost as much as they did in Darwin's day (Mayr & Provine, 1980). Severe communication problems, exacerbated by idiosyncratic terminology—the term “blending”, for example, referred to an intermediate *genotype* for the geneticists, whereas for students of whole-organisms (the naturalists) it referred to an intermediate *phenotype* (Mayr & Provine, 1980)—ensured that the camps remained at cross-purposes. Though

---

<sup>4</sup>The situation was actually more complicated than this, but as a simplification, this partition serves the purpose of this chapter (for a more detailed review, see Mayr & Provine, 1980).

both fields matured greatly in the early decades of the twentieth century, each severely under-appreciated the other, often misunderstanding the significance of the “opposition’s” discoveries.

The rediscovery of Mendelian genetics at the turn of the century initiated the formation of modern population genetics. Generally, it was believed that Mendelian genetics disagreed with Darwinism: rather than attributing much of evolution to natural selection, Mendelism elevated mutation to the primary source of evolution, considering adaptation as a result of “mutation-pressure”. Believing only in “hard” inheritance (genetic heredity only) the geneticists considered species in terms of *discrete* genetic systems. Resulting theoretical models led geneticists to predict that there was often not enough variation for natural selection to act. As such, new species were thought to originate from single macro-mutations—*saltationism*.

Contrastingly, the naturalists generally accepted gradual Darwinian evolution. Having secured a great wealth of knowledge of natural species, naturalists were well versed in the subtle (seemingly *continuous*) variations species displayed. Unlike the spatially independent discrete mathematical models that population geneticists (such as R. A. Fisher and J. B. S. Haldane) tended to use, naturalists understood the key role spatiality (in particular geographic isolation) played in the variation exhibited by natural species. Realising that natural populations contain enough variation for natural selection to occur, naturalists placed much less emphasis upon the role of macro-mutations in speciation. Unconstrained by Mendelian genetics, naturalists often utilised theories of “soft” inheritance (both genetic *and* non-genetic heredity), including the neo-Lamarckian evolution of acquired characteristics.

Given such markedly different paradigms it is surprising that, in the years between 1936-1947, a synthesis between Mendelian genetics and natural selection occurred. Marrying population genetics with natural selection, it was possible to finally understand how continuous changes in phenotypic traits can result from a discrete genetic system. The resultant neo-Darwinist paradigm buried, once and for all, the theories of soft inheritance and saltationism.<sup>5</sup> Many of the questions Darwin raised nearly a century before had finally been answered. In conjunction, the powerful mathematical tools developed by the geneticists could be applied to the wealth of knowledge the naturalists had gathered.

Laying the ghosts of the nineteenth century to rest, evolutionary biology was able to rapidly progress, flourishing in the unity of knowledge the formally disparate groups

---

<sup>5</sup>This is not strictly true. One may consider that a resurgence in these theories has occurred since the evolutionary synthesis. Whilst the theory of punctuated equilibria (Eldredge & Gould, 1973; Gould & Eldredge, 1977) can be *loosely* equated with saltationism, the theory of memes (Dawkins, 1976) is a strong contender for soft inheritance.



could collate. Within a decade of the synthesis, Watson and Crick (1953) had discovered the double-helix of DNA and with it the fundamental mechanism of replication. The final piece of the neo-Darwinist jigsaw slotted into place:

It has not escaped our notice that the specific pairing we have postulated immediately suggests a possible copying mechanism for the genetic material (Watson & Crick, 1953).

### 1.3.5.2 New Techniques: A New Perspective

#### Inclusive Fitness

Exemplified by Wynne-Edwards (1962), group selectionism—considering a group or species to be the unit of selection rather than the individual—persisted in the decades immediately following the evolutionary synthesis. Group selection gained popularity through its ability to explain *apparently* altruistic behaviour in organisms. Why, for example, do social-insects, such as ants, forego their own lives to save a nest if not for the good of the species? Group selectionism assumes that a population will increase its chances of survival if it can adapt to act harmoniously with other species and the physical environment (Thompson, 1994). Thus, group selection presupposes that coevolutionary systems (such as host-parasite interactions) will always become more mutually beneficial and *less* antagonistic through time. This view severely stunts possible research avenues into the evolution of interspecific interactions.

However, group selectionism was forcibly attacked during the 1950s and 1960s as many argued that evolution occurred through natural selection among individuals, *not* groups. The nail in the coffin for group selection came when Hamilton (1964) produced his brilliant work on “inclusive fitness”—the fitness of an adult individual cannot be calculated merely by its offspring, but must take into account the likely number of other’s offspring that include a proportion of similar genes. Inclusive fitness thus explains altruism towards kin and quickly explains the “mystery” of self-sacrificing ants. However, as Grafen (1991) points out, the concept of inclusive fitness is very complicated and often misunderstood:

... the scope of inclusive fitness covers all interactions in which the genotype of one individual affects the fitness of conspecifics. The special role of relatives is a powerful result of the theory, not a restricting assumption (Grafen, 1991, p9)

Inclusive fitness, and the death (or, at least, “exile”) of group-selectionism, meant that all interspecific interactions were no longer expected to result in mutualism.

These views opened the way for studying how different ecological conditions favoured different degrees of specialization and different outcomes of coevolution (Thompson, 1994, p48).

### Game Theory

Originally developed as a modelling tool for economic behaviour, *game theory*—the use of quantification methods to study competitive situations—enabled, for the first time, the analytical study of competitive games (von Neumann & Morgenstern, 1953). Unlike probability theory, which is constrained to analyse games of chance—the equivalent of “no-player” game theory—game theory assumes that rational agents engaged in a conflict of interest choose, at each turn, the most appropriate strategy. Outcomes are not random events but are determined by *rational* play.

By treating rationality as stability, and self-interest as Darwinian fitness, game theory was first applied to evolutionary biology in the 1960s. Paradoxically, though game theory was first developed as a tool for economics, it has proven more applicable to modern evolutionary biology (Maynard Smith, 1982, Preface).<sup>6</sup> Game theory enables evolutionary biologists to predict which, of a given set of strategies, will perform well within a population. Further, one can observe whether such a strategy, once used by the majority, can be *invaded* (outperformed) by some opposing strategy and whether such strategy invasions will cycle periodically, fluctuate chaotically or settle at an equilibrium. Such an equilibrium—a population converged upon a non-invadable strategy—is termed an Evolutionarily Stable Strategy (ESS; Maynard Smith & Price, 1973).

Over the last four decades, game theory has revolutionised evolutionary biology, finally permitting the formal analysis of evolutionary competition. In conjunction with population genetics, game theory—inherently considering the interaction between opposing strategies—offers a new technique for studying coevolution (e.g., Maynard Smith, 1982; Sih, 1998, Chapter 10).<sup>7</sup> Slatkin (1983) describes how this can be done:

---

<sup>6</sup>The reason for this is game theory’s assumption of rationality. Economic agents often *do not* act rationally—they are prone to uninformed decisions, biases and prejudice, they may be unprepared to take risks, or to follow whims, trends and fashions. However, by replacing rationality with population dynamics and stability (Maynard Smith, 1982), the assumption is no longer as fallible—Darwinian fitness cannot be “deceived” as easily as human rationality. However, as Gould and Lewontin (1979) argue, embryological and morphological constraints *can* lead to inherent biases in evolutionary dynamics—the open question is “to what extent?” In response, the evolutionary game theorist can return to the initial assumptions of the game, where each possible strategy is explicitly defined as equally achievable, thus side-stepping the anti-adaptationist argument.

<sup>7</sup>However, it should be recognised that game theory and population genetics tend to consider *fixed* equilibria, often overlooking the coevolutionary dynamics of continuous change (Maynard Smith, 1982, Introduction). Yet, game theory *can* handle cyclic trajectories if additional analytic techniques are applied (e.g., Maynard Smith, 1982, appendix J).

Game theory can be used to understand models of coevolving systems by envisioning each species as a “player” that is trying to maximize its gain through certain kinds of interactions under the assumption that the other species is trying to do the same (Slatkin & Maynard Smith, 1979). For example, the coevolution of two competing species can be understood in terms of the strategy of generalization leading to the use of more resources but causing more intense competition, as contrasted with a strategy of specialization leading to the use of fewer resources but also leading to less intense competition (Lawlor & Maynard Smith, 1976).

### Computer Models

The development and growth of the digital computer after World War II heralded a new age in evolutionary biology. As computational power slowly trickled into the hands of biologists, there emerged a new methodology for studying evolution. For the first time, it became possible to simulate evolving populations and observe interactions on the level of individual agents. Taking into account detailed logistical factors—such as population structure, migration or generation time—much more intricate models of interactions between both individuals and populations became available.

When modelling host-parasite coevolution, for example, it is necessary to capture not only the genetics of both populations, but also the spatial interplay of epidemiology. Achieving this mathematically, however, is fraught with technical complications. As such, mathematical models of host-parasite interactions tend to capture only one aspect of the system. Generally, these include parasite-host genetics with no epidemiology, host genetics and frequency dependent fitnesses, or population models with no genetics (May & Anderson, 1983).

Unlike population genetics and game theory, computer models “trivially” reproduce complex population interactions and subsequent coevolutionary dynamics. However, they are unfortunately much more opaque to analysis. The benefits of adding complexity with ease is countered by the fact that results gleaned from such complex simulations may be analytically intractable.

Throughout the 1960s, several attempts were made to harness the essence of evolution computationally. These included evolution strategies (Rechenberg, 1973), evolutionary programming (Fogel, Owens, & Walsh, 1966) and genetic algorithms (Holland, 1975).<sup>8</sup> Each used a very simplified analogy of evolution, including a population of individuals consisting of very simple genetic code, some measure of assessment, and a subsequent

---

<sup>8</sup>See chapter 2 for a review of evolutionary computation.

process of reproduction and mutation.

Implemented as a (non-evolutionary) computer simulation, the model of animal conflict by Maynard Smith and Price (1973) was one of the first compelling game theory models in evolutionary biology. Modelling conflict as a series of pairwise confrontations between two opponents, Maynard Smith and Price (1973) allowed adversaries a choice of three available actions each move; a conventional threat display, a potentially dangerous attack, or a retreat. Five opponent strategies—each a set of probabilities governing the choice of action—were defined. Each strategy played itself and the other four over a series of 2000 moves, with results demonstrating that “limited war” strategies (those including conventional displays) were evolutionarily stable, whilst “total war” strategies (those without conventional displays) were not. However, in a little cited letter to *Nature*, Gale and Eaves (1975) point out that an alternative ESS exists in the form of a frequency-dependent mixture of total war strategies. Using (perhaps) the first rudimentary evolutionary simulation model, Gale and Eaves (1975) demonstrate that, despite their absence as an equilibrium end-point in game theory, total war equilibria *will* be reached from certain initial conditions. This demonstrates a potential advantage of evolutionary simulation models; such techniques are able to consider evolutionary trajectories unavailable to game theory.

Since the evolutionary synthesis, the necessary tools to study coevolution have gradually developed. Computers afford intricate simulation models, game theory allows formal analysis of competition, and the principle of inclusive fitness offers a range of non-mutualistic coevolutionary outcomes. Marrying such techniques with the post-synthesis advances in genetics and ecology, evolutionary biology finally reached a position to tackle the fundamental questions of coevolution.

### 1.3.5.3 The Coining of Coevolution

In 1958, the first mathematical model of coevolution was published (Mode, 1958). Demonstrated by Flor (1955) to exist in interactions between flax and flax rust, Mode (1958) utilised the concept of gene-for-gene interaction, matching each gene for host resistance with an explicit gene for parasite virulence (see Thompson, 1994, p52-3).

However, it was not until 1964 that the term “coevolution” was first defined:<sup>9</sup>

---

<sup>9</sup>Although (Mode, 1958) included the term “co-evolution” in the title of his paper, Futuyma and Slatkin (1983) suggest that Ehrlich and Raven (1964) coined the term “coevolution”. Presumably, their reasoning follows that Ehrlich and Raven explicitly define “coevolution” in terms of interspecific interactions, rather than implicitly referring to “co-evolution” as some sort of ecological evolution or multiple evolution in a *generic* sense. As Ehrlich and Raven (1964) suggest, many studies of community evolution tend to be narrow in scope and ignore the *reciprocal* aspect of interactions, often considering one population as

... coevolution is the examination of patterns of interaction between two major groups of organisms with a close and evident ecological relationship, such as plants and herbivores (Ehrlich & Raven, 1964).

During their seminal work on the antagonistic coevolution between butterflies and plants, Ehrlich and Raven (1964) conclude, on the basis of an extensive survey of patterns affecting food plant choice, that the phylogeny of parasites (butterfly larvae) tracks the phylogeny of hosts (plants) during the continual evolution of defensive and counter-defensive measures. Within 3 years, Janzen's seminal work on the interdependency between the swollen-thorn acacias and their ant inhabitants paved the way for the study of mutualistic coevolution (Janzen, 1966). After the breakthrough of Ehrlich and Raven (1964) and then Janzen (1966), the late 1960s and early 1970s witnessed a dramatic increase in the attention evolutionary biologists bestowed upon coevolution, both antagonistic and mutualistic.

Perhaps the most significant publication of the early 1970s was van Valen's (1973) *A New Evolutionary Law*, in which he introduced his radical Red Queen hypothesis.<sup>10</sup> van Valen made an analogy between the Red Queen and biological coevolution after discovering a surprising trend concerning the probability of extinction of species. After analysing huge data sets collected across a wide range of biological taxa, van Valen noticed that, counter to intuition, "all groups for which data exist go extinct at a rate that is constant for a given group". Assuming that species evolve in a relatively static environment, one would expect beneficial adaptations to accumulate over evolutionary time, enabling progressive generations to be better equipped at defending against extinction than their ancestors. To explain his findings, van Valen proposed a new evolutionary law of extinction, with the Red Queen hypothesis as its central tenet: "biotic forces provide the basis for self-driving ... perpetual motion of the effective environment". Any beneficial adaptation by a particular species is, by definition, detrimental to other species inhabiting the same effective environment. Coevolutionary forces will, in turn, select for specific counter-adaptations in these species. In this way, adaptive advantage is continually eroded. In Red Queen fashion, species are truly *running to stand still*.

The Red Queen hypothesis was not only the first instance of an explicitly "general" coevolutionary theory, but also hinted at coevolution's potential to be one of the predominant forces driving evolution in nature. One may expect Red Queen dynamics to drive a run-away process of continual counter-adaptation—a coevolutionary *arms race*. Such

---

invariate.

<sup>10</sup>Appearing in Lewis Carroll's *Through The Looking Glass*, the Red Queen is a living chess piece that continually runs to stand still. No matter how fast she moves, the surrounding landscape always keeps up with her.

arms races are considered a profound force driving evolutionary adaptation in the natural world, and have sometimes been characterised as a source of strong selection for novel adaptations capable of accelerating evolutionary progress (Dawkins & Krebs, 1979).<sup>11</sup>

Throughout the 1970s, coevolution began to develop as a discipline in its own right, aided by the first textbook to dedicate an entire chapter to coevolution (Futuyma, 1979) and culminating in the appearance of entire works solely concerned with coevolution (Thompson, 1982; Futuyma & Slatkin, 1983). After two centuries of dabbling with interspecific interactions, evolutionary biologists finally had a firm foundation with which to study coevolution. However, the rise in emphasis upon coevolution led to the worrying situation in which coevolution had become such a pervasive concept that it was in danger of rendering itself moot. As a backlash to this (discussed in detail in Chapter 3) attempts were made to restrict coevolution by enforcing strict definitions.

## 1.4 Summary

Based upon the concept of interspecific interactions, biologists from the time of Darwin and beyond have entertained a “coevolutionary perspective” on the natural world. However, it was not until after the evolutionary synthesis that the emergence of new techniques facilitated the establishment of coevolution as an independent discipline. Yet despite the availability of new tools and a greater understanding of the natural world, due to their inherently counter-intuitive nature, coevolutionary systems will undoubtedly continue to elude the understanding of evolutionary biologists for some time.

In the following chapter we will see that evolutionary computation has, perhaps surprisingly, failed to take a coevolutionary perspective until very recently. Given that evolutionary computation was developed during the pinnacle expansion period of coevolution, it is interesting that interspecific interactions have largely been ignored.

---

<sup>11</sup>The concept of coevolution driving an evolutionary arms-race is central to coevolutionary computation in general, and this thesis in particular. See chapter 3 for a more detailed discussion.

# Chapter 2

## Review of Evolutionary Computation

---

Facilitated by the growth and accessibility of digital computers, the discipline of evolutionary computation (EC)—applying biologically inspired computational analogues to complex search problems, or to biological modelling—emerged during the 1960s. However, though interest in coevolution flourished amongst evolutionary biologists during this period (refer to chapter 1), interspecific interactions were mysteriously absent during the formative years of EC. Indeed, in marked contrast to evolutionary biology, only in the last decade or so has a “coevolutionary perspective” entered evolutionary computation.

Since evolutionary computation plays such an integral role throughout this thesis, this chapter endeavours to outline the fundamentals of EC in order to provide the reader with sufficient background knowledge. Specifically, the genetic algorithm (GA)—a tool used extensively throughout the experimental chapters 6 and 7—is introduced in detail.

### 2.1 Artificial Evolution

It is generally accepted by contemporary evolutionists that evolution *per se* does not require biological organisms in order to occur. As Lewontin (1970) suggests, the process of evolution requires three necessary (and sufficient) properties; variation, differential fitness, and heritability. Given that such requirements are fulfilled, a population is capable of evolving whether biological or not, thus allowing the possibility of “artificial” evolution—the evolution of non-biological entities.

Lewontin (1970) states that Darwin's scheme embodies three principles:

1. Different individuals in a population have different morphologies, physiologies, and behaviours (phenotypic variation).
2. Different phenotypes have different rates of survival and reproduction in different environments (differential fitness).
3. There is a correlation between parents and offspring in the contribution of each to future generations (fitness is heritable).

[These] embody the principle of evolution by natural selection. While they hold, a population will undergo evolutionary change (Lewontin, 1970).

In the form of an evolutionary algorithm (EA)—a collective term for the class of algorithms incorporating, amongst others, the genetic algorithm (GA)—digital computers afford the necessary framework with which to instantiate artificial evolution. Though EAs vary dramatically, the underlying conceptualisation of biological evolution is shared by all; whether used for evolutionary simulation modelling or for parameter optimisation, evolutionary algorithms are inherently population-based search techniques that incorporate some form of heredity, variation and selection.<sup>1</sup>

Individuals within a population are defined using a simple encoding scheme (analogous to DNA) before being assessed in order to assign each a “fitness” score, roughly approximating adaptive success. With success breeding success, individuals reproduce in proportion to their fitness. Aside from the occasional mutation, children are replicas of parents. In this manner, EAs utilise Darwin's principle of “survival of the fittest”, attempting to drive the population to become ever more successful at the desired task.

Following World War II, as computers began to pervade evolutionary biology, many forms of evolutionary algorithm sprang into existence, eventually resulting in three recognised “schools” that persist, to varying degrees, today. In Germany, Rechenberg (1973) developed the evolution strategy, or *Evolutionsstrategie* (ES), whilst around the same time two alternative approaches, evolutionary programming (EP, Fogel et al., 1966) and genetic algorithms (GAs, Holland, 1975), emerged in the USA. Though other forms of evolutionary algorithm were developed (e.g., Reed, Toombs, & Barricelli, 1967), they were not widely adopted.

---

<sup>1</sup>This is perhaps misleading since evolutionary algorithms occasionally utilise a population size of one; a “(1+1)EA”. Each generation, a mutant offspring is created from the lone population member and the two are compared. The winner is retained for the following generation whilst the loser is discarded. Commonly referred to as “hill-climbers”, such algorithms exhibit very different dynamics to multi-individual population EAs.



Although [ESs, EP and GAs] share an identical meta-model, each one emphasises different features as being most important to a successfully modelled evolutionary process (Bäck, 1996).

Apart from the variety of encodings used, the most salient differences between the algorithms is the emphasis that each places upon the genetic operators (mutation and recombination) and population size.<sup>2</sup> Whereas Holland's GA relies upon genetic recombination and mutation in multi-individual populations, evolutionary programming generally utilises only mutation, whilst evolution strategies historically incorporate a (1+1) population hill-climber. As Mitchell (1996) suggests: "Holland was the first [to] attempt to put computational evolution on a firm theoretical footing". Unlike EP and ESs, GAs were not designed purely for function optimisation, but rather as a way to accurately model the evolutionary process. As such, GAs more realistically capture the essence of evolution than competing algorithms, perhaps influencing their popularity over recent decades (see section 2.3 for a discussion of realism).

It is important to highlight the fact that, over the past three decades, the strict classifications discussed here have not remained steadfast. Many algorithms now form some sort of hybrid inspired by several schools of thought. Yet, the majority of work covered in this thesis fits most comfortably into the GA class. For this reason, and because of the pervasiveness of GAs in the EC community, the following section introduces GAs in detail.

## 2.2 Genetic Algorithms (GAs)

GAs often vary radically from the original design of Holland. Indeed, only ten years after Holland (1975), De Jong (1985) commented upon the metamorphosis of GAs, suggesting that: "today, researchers often use the term "genetic algorithm" to describe something far from the original conception of Holland". Consequently, there is no rigorous GA definition. Yet, as Mitchell (1996) suggests, most GAs (including Holland's) share some common elements: populations of individual genotypes, selection according to fitness, crossover to produce new offspring, and random mutation of offspring.

This section introduces a simple generic GA containing only Mitchell's "core" components. Hopefully, this should be enough to elucidate the standard GA approach without

---

<sup>2</sup>Traditionally, each algorithm is associated with a particular encoding scheme; ESs with real-valued  $n$ -dimensional vectors, EP with finite state machines and GAs (at least initially) with binary strings. Such associations, however, are not ubiquitous. In particular, having the ability to accommodate a wide variety of encodings, GAs frequently utilise alternatives to binary string encodings in contemporary work.

becoming entrenched in, possibly idiosyncratic, details. Throughout the remainder of this section, Mitchell (1996, Chapter 1) is utilised as a basis.

### 2.2.1 A Simple Generational GA

A simple GA often follows the following plan:

1. Randomly generate an initial population of  $n$  individuals, each consisting of a genotype of length  $l$ . Set generation counter  $g = 0$ .
2. Using a pre-defined function, calculate the fitness  $f(x)$  of each individual,  $x$ , in the population.
3. Repeat the following until  $n$  offspring have been created.
  - (a) Select, with replacement, two “parent” individuals, where the probability of selection is a monotonically increasing function of fitness (refer to section 2.2.4).
  - (b) With probability  $p_c$ , apply crossover to the parents’ genotypes (refer to section 2.2.3.2).
  - (c)  $\forall i \leq l$ , mutate new offspring at location  $l_i$  with probability  $p_m$  (refer to section 2.2.3.1).
4. Replace the current population with the new population of offspring. Increment counter  $g$ .
5. *If* ( $g < g_{max}$  & solution not found) goto step 2, *else* end.

Many GA applications follow this simple procedure. However, the success of an algorithm depends upon the details of the particular procedure chosen. Such details include the population size  $n$ , the number of generations  $g_{max}$ , and the probabilities of mutation  $p_m$  and crossover  $p_c$ . Also, the modes of variation (see section 2.2.3), assessment and selection (section 2.2.4) play an important role. Often, some effort is needed to tweak these aspects in order to achieve sufficient success, either in modelling or optimisation.

A simulation model can be classified successful if it accurately reflects an abstraction of the system it is modelling; such simulations can then be used to develop new hypotheses about the natural system of interest. In contrast, a successful optimisation algorithm is one that generally discovers good solutions to a required problem in a “reasonable” amount of time. For optimisation, the realism of an algorithm is largely irrelevant (see

section 2.3). The following section details a worked example of a problem solved by a GA.

## 2.2.2 Worked Example

### The Problem

Using the 53-letter alphabet  $\Sigma = \{a, b, \dots, z, \text{ }, A, B, \dots, Z\}$ , containing the 26 lower-case letters, 26 upper-case letters and space character, find the string of 19 characters equal to: “Rules of Engagement”.<sup>3</sup> Although this problem may at first appear trivial, there are  $19^{53}$  possible combinations of 19-character strings. Thus, using an exhaustive search would require an inordinate amount of time ( $2.97 \times 10^{67}$  iterations, on average).

### Set Up

Rather than use a binary encoding for the genetic algorithm it seems sensible, given the nature of the problem, to represent each individual as a 19-character string containing letters drawn from the alphabet  $\Sigma$ . To begin, set generation counter to 0 and initialise the population of  $n = 100$  individuals *randomly*.

Each generation, individuals are assessed by calculating the number of loci which match the solution string. Thus,  $g_1 = \{\text{Ruses of EngKgmmBnt}\}$  receives a fitness  $f(g_1) = 15$ , whilst  $g_2 = \{\text{Rules of Engagement}\}$  receives the maximum fitness  $f(g_2) = 19$ . If the solution string has not been found and the number of generations is less than  $g_{max} = 1000$ , then the population is allowed to reproduce. Using fitness proportionate selection (biased selection favouring high scoring individuals, see section 2.2.4), offspring are created using 1-point sexual crossover (crossover location chosen randomly, see section 2.2.3.2) and point mutation (section 2.2.3.1).

Let  $p_1 = \{\text{Rules of EngaJement}\}$  (fitness 17), and  $p_2 = \{\text{NulesYof Engagement}\}$  (fitness 17), be chosen to reproduce. The crossover point is chosen at locus 12. Hence,  $\{\text{Rules of Eng}\}$  is copied from parent  $p_1$  and  $\{\text{agement}\}$  is copied from parent  $p_2$ . Occurring at each locus with probability  $p_m = 0.02$ , mutation affects *only* locus 17, flipping allele “e” to a randomly chosen alternative, “B”, thus resulting in the final offspring  $h = \{\text{Rules of EngagemBnt}\}$  (fitness 18).

Once 100 offspring are generated the parent population is replaced. The generation counter is incremented and the assessment phase is repeated.

### Results

Figure 2.1 displays the result of one evolutionary run. Each time a new individual is discovered, capable of outperforming—containing more correct loci—all previous indi-

<sup>3</sup>This example is functionally equivalent to Dawkins’ (1986) evolution of the sentence “Methinks it is a weasel”.

viduals, its genotype, fitness score and generation of first discovery are output. After just 11 generations a solution appears with more than half of all loci correct. The run completes when an optimal solution evolves at generation 44; having searched a total of 4500 individuals: very few in comparison to the expected number of iterations a blind search is expected to take, on average.

Gen	Score	Best Genotype (So Far)
000	4	RVFIYLrgJbnICgDmBKT
002	5	RVFIYLrgJbnICgDmBKt
004	6	jqldbDoYnbnICgDmBKt
005	7	LursMofuEIDaVTRfYt
006	8	Ryue LEezEngCgDmBnS
007	9	RVFxoLofJbngxgemqqt
011	11	RulesmoQ EXCgDmsnZ
013	12	RulesYCh EnhCgDmwnt
017	14	RulesYof EngCgeYsnZ
028	15	Ruses of EngKgmmBnt
029	16	NulesYof EngaJement
030	17	RulesYof EngaJement
042	18	Rules of EngagemBnt
044	19	Rules of Engagement

Figure 2.1: Results from a simple GA. Initially randomised, individuals are encouraged to evolve towards the *solution* genotype “Rules of Engagement”. Each new best individual is output. After just 44 generations an optimal solution is discovered.

## 2.2.3 Modes of Variation

### 2.2.3.1 Mutation

The second law of thermodynamics states that any copying system will suffer occasional errors. Thus, in biological systems, mutation is an unavoidable consequence of the universe we inhabit.<sup>4</sup> However, though mutations are often considered deleterious, the evolutionary process is fundamentally reliant upon them as a source of variation.

<sup>4</sup>Error-checking mechanisms, commonly found in biology, vastly reduce the probability of mutation. However, such mechanisms are energy expensive. Thus, no copying mechanism is perfectly error free.

The rate of mutation,  $p_m$ , is an important factor in evolution. If  $p_m$  is too high, then beneficial adaptations are likely to be lost; for each beneficial gene, there is a probability  $p_m$  that it will be replaced by a random, often deleterious, alternative. However, if  $p_m$  is too low then there is little chance of any gene mutating, whether beneficial or not. As such, relatively few *new* alleles are able to enter the “gene-pool”, thus reducing the level of variation within the population. As alleles are shuffled through sexual recombination (or reproduced unchanged in asexual systems), it is likely that the population will begin to converge upon one genetic clone. Although neither situation is acceptable for “successful” evolution, whilst *some* variation is necessary, less is usually better. Accordingly, the rate of mutation is often kept low. Following biological systems as a guiding heuristic, GAs often utilise a mutation rate of approximately one expected mutation per offspring,  $p_m = \frac{1}{l}$ .

The most frequently used mutation in the GA community is the point mutation. In general, given a discrete genetic alphabet,  $\Sigma$ , containing  $\sigma$  characters, when mutation occurs at a particular locus, the current allele is exchanged for a randomly selected allele  $\alpha \in \Sigma$ . If an ordering is imposed upon  $\Sigma$ , mutations may move either *up* or *down* the alphabet. Thus, if  $\Sigma = \{0 < 1 < 2 < \dots < 9\}$  then allele 3, say, may mutate to either 2 or 4. If a GA is encoded with continuous real-values, mutations generally occur by adding a random value (drawn from a gaussian distribution with mean 0) to the current value. However, susceptible to bias, such mutations may drive the evolving population into particular regions of space (see, e.g., Bullock, 1999, 2001). More unusual encodings (such as finite state machines) require more specialised forms of mutation (see, for example, Koza, 1993; Gruau, 1995).

Inspired by biological systems, more complicated forms of mutation may also be used within a GA. Translation mutations copy a chunk of genotype from one location to another, whilst others, such as Holland’s (1975) inversion operator, swap the characters found at two, or more, locations. Finally, given a variable-length encoding, mutation may either increase or decrease the length of genotype.

### 2.2.3.2 Recombination

Genetic recombination, or crossover, is a mechanism for producing new offspring genotypes through the combination of two parent genotypes in a sexual GA population. Recombination accelerates the transfer of genetic information throughout a population and reduces the similarity between a parent and its offspring; in biology, one theory (of many) suggests that sexual recombination may have evolved to combat the effects of parasitism (as offspring are likely to enter an environment in which antagonistic parasites have

specifically adapted to their parents, genetically variable offspring that are not identical to parents may have a higher resistance to parasites, Hamilton, 1980), whilst another suggests that a changing environment is responsible (Waxman & Peck, 1999).

Crossover generally comes in three “flavours”. To illustrate these, let us assume that parents  $p_1 = \{111111\}$  and  $p_2 = \{000000\}$ , each of length  $l = 6$  are selected for recombination.

**Single-point.** Randomly choose crossover position  $c$  such that  $0 < c \leq l = 6$ . Form new (offspring) genotype by copying the genotype of  $p_1$  from loci 1 to  $c$  inclusive, then subsequent loci copied from  $p_2$ . For example, if  $c = 3$ , then offspring =  $\{111000\}$ .

**Multi-point.** Randomly choose  $n$  crossover positions such that  $0 < c_1 < c_i < c_n \leq l = 6$ , then copy  $p_1$  from loci 1 to  $c_1$  inclusive,  $p_2$  from  $c_1 + 1$  to  $c_2$  inclusive,  $p_1$  from  $c_2 + 1$  to  $c_3$  inclusive, etc. For example, suppose  $n = 2$  with  $c_1 = 2$  and  $c_2 = 5$ , then offspring =  $\{110001\}$ .

**Uniform.** For each locus, chose with equal probability whether the copied character will be from  $p_1$  or  $p_2$ . For example, assume loci 2, 5, 6 copied from  $p_1$  and loci 1, 3, 4 copied from  $p_2$ , then offspring =  $\{010011\}$ .

Importantly, it should be noticed that recombination is a process for mixing genetic information within a population; it *is not* a mechanism for creating *new* genetic material. Without mutation, a sexual population will eventually become fully converged upon one genotype. At such times recombination becomes redundant (see, for example, Harvey, 1993).

## 2.2.4 Assessment and Selection

Each generation, it is necessary to assess the “fitness” or “success” of each individual within the GA population in order to determine which should contribute more, or fewer, offspring to the following generation. Fitness assessment must determine which individual “solutions” are better than others in terms of the designated problem at hand. As such, the assessment function chosen is strongly linked to each particular problem and is very task-specific. However, some important considerations always apply. Assessment should not demand excessive computation and should discriminate between individuals throughout the entire evolutionary run, irrespective of whether the population consists of relatively poor, moderate or high fitness individuals.

Following assessment, individuals are selected to reproduce. Based upon the fitness assigned to each individual, the following methods of selection are commonly used within the GA community.

**Fitness proportionate selection.** Originally used by Holland (1975), fitness proportionate selection directly maps an individual's fitness with the number of offspring that individual will be expected to produce. In similar fashion to standard evolutionary biology models, Holland's expected number of offspring is calculated by dividing the fitness of an individual by the average fitness within the population. Thus if  $f(x) = 3$  and average population fitness  $f(\bar{x}) = 1.5$ , then individual  $x$  would be expected, on average, to produce 2 offspring. This procedure maintains constant population size.<sup>5</sup>

Fitness proportionate selection is most often implemented using the *roulette wheel* approach. Each individual is allocated a slice of the wheel, with size proportional to fitness. The wheel is then spun and whichever slice the pointer lands upon, that individual is chosen to breed. This is repeated  $n$  times, where  $n$  is the size of the population.

**Rank-based selection.** Rank-based selection (e.g., see Baker, 1985) was introduced in order to avoid premature convergence within a population. As individuals become similar, the absolute difference in fitness falls, leaving proportionate selection struggling to discriminate. If the best individuals only expect 0.001 more offspring than the average individual, then clearly, being best no longer holds significance (assuming population size is not many orders of magnitude greater than 1000).

Rank-based selection avoids the problem of convergence faced by fitness proportionate selection by calculating the expected number of offspring through an individual's fitness ranking within a population, rather than absolute fitness score. Thus, no matter what the difference between the best individual and the worst, the expected number of offspring of each never varies. In this way, selection pressure (the advantage gained from being a fit member of the population) remains constant throughout a run.

---

<sup>5</sup>Let  $n$  = current population size,  $n'$  = new population size,  $f(\bar{x})$  = mean population fitness and  $\sum f(x)$  = sum of population fitness, then

$$n' = \sum \frac{f(x)}{f(\bar{x})} = \frac{\sum f(x)}{f(\bar{x})} = \frac{\sum f(x)}{\frac{\sum f(x)}{n}} = n$$

**Tournament Selection.** Avoiding the necessity to sort a population into ranked order, tournament selection is a computationally efficient way of implementing rank-based selection. Two individuals are chosen at random from the population. A random value  $r$  is chosen between 0 and 1. If  $r \leq k$ , where  $k = 0.75$  say, then the fittest individual is chosen to reproduce, else the less fit individual is chosen. Both are then replaced back into the original population.

Throughout this thesis, tournament selection will generally be used, however, the method employed varies somewhat from that described above. Tournaments containing  $t$  individuals are chosen, with the fittest individual *always* selected to reproduce ( $k = 1$ ). This approach forces much stronger selection pressure, allowing the fittest individuals more expected offspring, whilst weaker members of the population become less likely to reproduce.

**Elitism.** First used by De Jong (1975), elitism is a technique to be used in conjunction with either fitness proportionate, rank-based, or tournament selection. Designed to prevent the destruction of successful genotypes during reproduction, elitism ensures that the best  $q$  individuals of each generation are copied *unchanged* into the following generation. As a result, good adaptations can only be displaced by better adaptations; they cannot be lost by chance. Though elitism often performs well in optimisation, it is a very unnatural addition to the GA.

**Steady-State.** Genetic algorithms generally follow the *generational* model described in section 2.2.1. However, alternative steady-state models are not uncommon. Rather than have a fixed population size, with new individuals replacing old individuals each generation, a steady-state GA places new offspring into the same population as their parents. This is equivalent to having generational overlap, or *elitism*, the fraction of which has been termed the “generation gap” (De Jong, 1975). Steady-state models are sometimes preferred when modelling natural systems if such generational overlap is more realistic. Considering human populations, it is common for grandparents and grandchildren to co-exist. When modelling human social interactions, therefore, a steady-state approach would be the obvious choice. However, apart from the generation gap, steady-state selection comprises much the same features as generational selection.



## 2.3 Realism

Throughout this thesis, an implicit concern reoccurs; as an analogue of natural evolution, just how “realistic” is the standard GA? The relevance of this question, however, varies greatly depending upon a GA’s particular usage; whilst realism may, on occasion, be of utmost importance, at other times it becomes no more than a distraction. As a preemptive measure to alleviate possible confusion, this section attempts to address the issue of realism in GAs (applicable to EAs in general), specifically commenting upon why and when realism is important. In conjunction, some unrealistic, yet intrinsic, aspects of the standard GA are outlined.

Necessarily, successful scientific theories must, to some degree, “realistically” characterize nature. However, it is not necessary for a theory to accurately *copy* a natural system; indeed, through the introduction of simplifying assumptions, no theoretical model (by definition) entirely replicates reality. Determined by the encompassing scientific paradigm—underpinning which aspects of reality should be incorporated and which are unimportant—only the “pertinent” aspects of reality are reflected in a theory.

For example, consider Newton’s classical theory of gravitation. According to Newton, when modelling gravity, the pertinent aspects of a system are the relative mass and location of objects. An object’s colour, for instance, is irrelevant. As such, under the paradigm of Newtonian gravitation, a realistic model would accurately reflect the mass and location of objects whilst ignoring other extraneous variables such as colour, roughness, edibility, etc. The addition of colour, whilst more accurately copying nature, would not improve (when considering classical theory) the accuracy of the model in predicting future system states. In general, a scientific model should not attempt to copy nature *per se*, but should rather attempt to capture the theoretical ideas that have been derived from nature.

Evolutionary algorithms are generally utilised for one of two specific purposes; either the modelling of natural systems, or as an engineering tool for solving complex search problems (optimisation). Reflected in the details of an EA’s construction, each task requires a very different approach. Whilst modelling necessitates a certain amount of realism—determined by the parts of a system that theory predicts will affect evolution—optimisation yields no such constraint upon algorithm design.<sup>6</sup> Such antagonistic requirements have directed the development of EAs throughout their history.

Holland’s original aim for the genetic algorithm was to develop “a mathematical framework which makes it possible to extract and generalize critical factors of the biological processes” (Holland, 1992, original Preface). Holland clearly understood that

---

<sup>6</sup>Although optimisation is less constrained by theory, the underlying theoretical assumption that evolution is in some sense an “optimisation process” is implicit.

when modelling a biological system it is important that the simulation, where evolutionary theory deems relevant, accurately reflects the natural system. For instance, when modelling multiple demes within a population, the rate of migration should be similar to the biological migration rate. Unrealistic assumptions in a model may lead to artefactual results, thus rendering the simulation invalid.

Affording micro-interactions between agents as well as simulation complexity, repeated trials and control of variables, GAs offer a powerful research mechanism for evolutionary biologists. Though the biological community has been slow to embrace evolutionary simulation models, some notable successes have taken place, including pioneering work on the *Baldwin Effect* (Hinton & Nowlan, 1987), measuring evolutionary adaptations (Bedau, 1996), and modelling of sexual selection and ecosystems (see Mitchell, 1996).

Though Holland strove for a framework with which to model the evolutionary process, he was also aware of the GA's potential applicability to parameter optimisation (first realised by De Jong, 1975). Now widely applied, GAs have achieved successful implementation in a diverse set of problem domains, including evolving computer (Lisp) programs, evolving cellular automata to perform computations, predicting dynamical systems and protein structures, evolving neural networks, and evolving learning rules (see, e.g., Goldberg, 1989; Davis, 1991; Mitchell, 1996).

In contrast to evolutionary modelling, when applying evolutionary algorithms to problem solving, the importance of realism is subtly different. As the goal of optimisation is to produce adequate solutions to a problem in “reasonable time”, rather than accurately model a natural system, it is irrelevant whether an optimisation algorithm accurately reflects any natural system or not. However, it *is* important that EAs, when used for optimisation, realistically capture the fundamental essence of evolutionary theory.

Though realism may not be a prerequisite for optimisation, it is quite possible for natural systems to *inspire* the innovative design of EAs (clearly demonstrated by the fact that evolutionary algorithms are used for optimisation at all). As a source of novel ideas, nature plays an important role in artificial evolution for optimisation.

Yet, one must be careful not to fall into the trap of assuming that realistic algorithms will necessarily improve performance. Just because nature is a specific way does not imply that an abstraction of nature will be good for optimisation. Though nature may inspire, realistic additions to an algorithm should be discarded if no improvement in performance can be proven. To assume that the natural way is the best way—that because we see it around us it must be optimal—is to fall into the naturalistic fallacy (Moore, 1903). It is impossible to deduce a normative statement from a factual state; just because something

is a particular way does not mean that it *ought* to be that way. Hence, when applied to optimisation, an evolutionary algorithm should only be adapted to improve performance and not to increase realism for its own sake.

### 2.3.1 Fixed Population Size

In the standard generational GA, one underlying assumption is that of fixed population size; each generation, the current population is substituted for one of identical size. Computationally, this assumption makes perfect sense, removing the necessity to accommodate a fluctuating population whilst avoiding catastrophes, such as extinction events, which may prematurely terminate a run. However, by fixing population size, a genetic algorithm loses much realism; population dynamics play an integral role in evolutionary biology. As a result, some evolutionary phenomena, such as disengagement (or extinction), vary between artificial and natural systems.<sup>7</sup>

### 2.3.2 Genotype-Phenotype Mapping

In nature, there exists a clear separation between the genotype (the underlying genetic encoding) and phenotype (all the “effects” of a gene upon the world)<sup>8</sup> of biological organisms. The genetic structure of an organism largely determines phenotypic structure, however, embryological and environmental factors also play a key role, resulting in a highly non-linear mapping from genotype to phenotype.

During the evolutionary synthesis (see section 1.3.5.1) it became generally accepted that the influence between genotype and phenotype runs in one direction only. Whilst a change in genotype is likely to induce some (perhaps unpredictable) phenotypic trait, changes in phenotype *do not* affect the underlying genotype. As such, the genotype-phenotype mapping allows life-time adaptations and learning whilst preserving the genetic structure to be passed on *unchanged* to progeny. In conjunction, deleterious changes to the phenotype, such as the loss of a limb, are not passed to offspring.

In contrast to natural systems, the standard genetic algorithm generally blurs the relationship between genotype and phenotype. For simplicity, it is often assumed that one

---

<sup>7</sup>Such disparities between phenomena in natural and artificial systems is a recurring theme throughout this thesis.

<sup>8</sup>It is useful to consider Dawkins’ definition of the *extended phenotype*: “As always, ‘effect’ of a gene is understood as meaning in comparison with its alleles. The conventional phenotype is a special case in which the effects are regarded as being confined to the individual body in which the gene sits. In practice it is convenient to limit ‘extended phenotype’ to cases where the effects influence the survival chances of the gene, positively or negatively” (Dawkins, 1982).

directly encodes the other, i.e., genotype=phenotype.<sup>9</sup> However, such a simplification destroys the possibility of generative encodings, life-time adaptations and learning. Thus, not only is realism lost, but the opportunity to apply these powerful mechanisms to optimisation and problem solving is also removed.<sup>10</sup> To Holland's (1992, Chapter 3) credit, genotype-phenotype mappings are considered in his early work, with a discussion of epistatic interactions between alleles leading to non-linear changes in phenotypic traits.

### 2.3.3 Static Environment

The general consensus since the days of Lamarck is that the world is not static (see section 1.3.2). However, this has not been widely reproduced in genetic algorithms. Though populations constantly evolve within a genetic algorithm, the assessment function generally does not; irrespective of conspecifics, a particular individual will always achieve the same fitness score (assuming a deterministic fitness function). Given that a GA's assessment function equates to the "environment" of an evolving population, this is perhaps unrealistic. As natural populations evolve within an adapting environment, more realism could be achieved by utilising a dynamic assessment function, whereby the fitness of an individual (not to be confused with reproductive success) is free to *vary* through time.

One may argue that a static environment can be considered accurate when modelling a relatively short (in evolutionary terms) time period. However, given that a GA population often evolves from some random configuration to one that contains highly specialised individuals, the relative evolutionary timescale to achieve this "progress" in the biological world undermines this argument; by the time such evolution would have occurred, the world is likely to be very a different place.

Undeniably, by allowing a GA to evolve in a static environment, predictability and interpretability each become a much easier task. Given that Holland's initial intention was to set up a rigorous mathematical framework for understanding biological evolution, it is likely that simplicity was a primary factor in choosing the static model. However, it is disappointing and a little surprising that early GA protagonists did not consider the possibility of dynamic environments, especially given the surge in enthusiasm toward coevolution around the time of Holland's initial work (see section 1.3.5.3).

---

<sup>9</sup>This is not strictly true since a distinction between genotype and phenotype will always exist in some sense. Irrespective of the complexity of the genotype-phenotype mapping, it is always true that the genotype alone is acted upon by the genetic operators of recombination and mutation, whilst the phenotype is similarly subject to assessment and selection.

<sup>10</sup>Non-trivial genotype-phenotype mappings have appeared in more recent years, for example, in the evolution of artificial neural networks (Beer & Gallagher, 1992), or the generative encodings of Hornby and Pollack (2001). However, in no sense does *any* GA exhibit a mapping approaching the complexity of biological systems.

## 2.4 Summary

Developed by Holland (1975), the genetic algorithm is a versatile tool that can be used for both optimisation and biological modelling. Though Holland's GA was greatly inspired by biological evolution, little attempt was made to incorporate coevolutionary ideas for many years; a surprising fact considering the surge in interest in coevolution amongst evolutionary biologists during this period. In the following chapter, *coevolutionary* algorithms are treated in detail.

# Chapter 3

## Coevolutionary Computation

---

The coevolutionary perspective has profoundly affected the accepted understanding of the natural world (chapter 1). Further, having become a subsidiary of evolutionary computing, coevolutionary algorithms now present great potential as a tool for both problem-solving and biological-modelling (chapter 2). However, it is still not entirely understood what coevolution is—what *constitutes* coevolution?

Highlighting the difficulty of defining such a multi-faceted phenomenon, this chapter introduces the variety of guises in which coevolution appears. A novel definition of coevolution (to be used throughout this thesis) is then proposed. Using the concepts of “fitness propensity” and “fitness realization” (Mills & Beatty, 1994), a hard distinction between coevolution and *non*-coevolution is drawn. The expected advantages of utilising coevolutionary (rather than non-coevolutionary, or “standard evolutionary”) techniques for problem solving are discussed. A brief overview of historical successes illustrate coevolution’s potential. Finally, problems regarding coevolution’s idiosyncrasies are detailed: the chapter ends with a critique of the literature regarding previous attempts at combating these problems.

### 3.1 Background

A major transition in the development of GAs occurred at the beginning of the 1990s when a dynamic environment was first introduced in the form of a coevolutionary host-parasite

GA (Hillis, 1990). Having steadily risen in usage since Hillis' pioneering work, coevolutionary algorithms are now familiar practice in the field of evolutionary computation and offer two immediate advantages over standard evolutionary approaches.

Firstly, coevolution (such as the competitive framework used by Hillis) removes the problem of defining a fitness function with which to assess individuals. Let us consider the standard evolutionary approach of the worked example shown in section 2.2.2. Given the purposely simplistic nature of this example, the assessment function is trivial; the fitness of an individual equals the number of loci containing alleles that match that of the model sentence "Rules of Engagement". In more complicated scenarios, however, the definition of an assessment function can be tricky, if not impossible; how to operationally define "tennis ability" or "amusement", for example? Using competitive coevolution, it is only necessary to be able to assess which of two individuals can beat the other. Thus, setting up a simple pairwise competition between individuals suffices. In Hillis' model, the parasite lists (through competition with hosts) remove the necessity for defining an assessment function for host networks.

Secondly, coevolution allows the potential for "open-ended" evolution; that is, evolution that is not constrained to a fixed optimum, such as biological evolution. Unlike the worked example of section 2.2.2, or Hillis' sorting networks, open-ended evolution would be captured in a hypothetical domain whereby robot foxes and robot rabbits are coevolved, with foxes attempting to pursue evading rabbits. Given a suitably rich environment and strategy space that allows sophisticated behaviours to evolve, it would be theoretically possible for fox and rabbit strategies to continually become more sophisticated in a seemingly endless fashion. This exciting potential of coevolutionary computation is a significant advantage over standard evolution.

In general, coevolutionary algorithms can be loosely categorised into two flavours, *cooperative* and *competitive* (or antagonistic) coevolution. Despite the considerable discrepancy in methodology between the two, each offers a similar "advantage" over the standard evolutionary approach.

Cooperative coevolutionary algorithms are often utilised in situations where a problem can be naturally decomposed into sub-components. Individuals represent such sub-components and are assessed in a series of *collaborations* with other individuals in order to form complete solutions (for example, Potter & De Jong, 2000; Watson & Pollack, 2000; Wiegand, Liles, & De Jong, 2001). Cooperative coevolutionary algorithms have had success in a variety of domains, for example, manufacturing scheduling (Husbands & Mill, 1991), function optimisation (Potter & De Jong, 1994), designing artificial neural networks (Potter & De Jong, 1995) and room painting (Puppala, Sen, & Gordin, 1998).

Competitive coevolution either occurs within one population engaged in self-play, or between multiple populations. Individuals represent *complete* solutions that are gradually refined throughout an evolutionary run. Single population competitive coevolution has been successfully applied to the Iterated Prisoner's Dilemma (Axelrod, 1984; Lindgren & Nordahl, 1994), pursuit and evasion (Reynolds, 1994), and to finding robust game strategies in, for example, Tic-Tac-Toe (Angeline & Pollack, 1993), backgammon (Pollack, Blair, & Land, 1996) and Texas Hold'em Poker (Noble & Watson, 2001).

This thesis is primarily concerned with competitive coevolution between two populations; the simplest and most commonly used  $N$ -population competitive algorithm, with inter-population assessment. Such competitive coevolutionary algorithms have been successfully applied to discovering minimal-length sorting networks (Hillis, 1990; Juillé, 1995), finding CA rules to solve the density classification task (Juillé & Pollack, 1998b, 1998a), designing artificial neural networks for robot control (Floreano & Nolfi, 1997a), pursuit and evasion (Cliff & Miller, 1995), and the domains of 3-D Tic-Tac-Toe and Nim (Rosin & Belew, 1997). For a more comprehensive list of successful coevolutionary applications refer to section 3.3.

## 3.2 Defining Coevolution

Often, the significance of a rigorous definition of coevolution is overlooked. When exclusively utilising coevolutionary algorithms for problem solving, a rigorous definition may be largely perceived as irrelevant—if it works, it works, no matter what you call it. Thus, artificial coevolution is sometimes considered more a subset of machine learning than a realistic extension of biological coevolution.<sup>1</sup> However, whilst this approach may not inhibit problem solving, it is unlikely to produce insight into the general principles of coevolution.

Without a strict definition, it becomes difficult to distinguish between what *is* and what *is not* coevolution. Potentially, this can lead to a situation in which anything evolutionary can be considered coevolutionary in some sense. Consequently, the term becomes moot. As a preventative measure, this section endeavours to formulate a definition of coevolution that is consistent with *both* biological *and* artificial systems, as generally described.

---

<sup>1</sup>Following his talk entitled “Beyond competition (& co-operation): A new principle for progress”—presented on September 17, 2003, at the European Conference on Artificial Life, Dortmund, Germany—Jordan Pollack appeared to support a similar view. When asked whether the work performed under his supervision in the DEMO lab (Brandeis University, MA) should fall under the classification of coevolution, Pollack replied that the biological definition is probably not entirely consistent with his work, concluding that it may be more appropriately paralleled by the machine learning approach of self-play.



It is hoped that a rigorous definition will allow greater knowledge transfer between evolutionary biology and evolutionary computing: a potential benefit to both fields.

### 3.2.1 The Standard Definition of Coevolution

Towards the end of the 1970s it became clear that without a universally accepted definition, coevolution was in danger of losing practical significance. Having been manipulated by various authors to fit a multitude of disparate phenomena, the term “coevolution” began to encompass virtually all biological interactions. As such, the distinction between coevolution and non-coevolution (or evolution in general) gradually disappeared.

The original, broad definition of the word *coevolution*, introduced by Ehrlich and Raven...embraces all evolution resulting from biological interactions (Stanley, Van Valkenburgh, & Steneck, 1983).

... adaptation to an effectively constant feature of the biotic environment does not differ from adaptation to a constant feature of the abiotic environment. Coevolution, too broadly defined, becomes equivalent to evolution (Futuyma & Slatkin, 1983, introduction).

As a direct backlash, Janzen (1980) produced a short paper entitled *When is it coevolution?* Explicitly designed to construct a strict definition, Janzen dismissed the common misuse of coevolution as synonymous with “interaction”, “symbiosis” and “mutualism”. Today, Janzen’s definition (in some form) is readily accepted by the majority of biologists.

‘Coevolution’ may be usefully defined as an evolutionary change in a trait of the individuals in one population in response to a trait of the individuals of a second population, followed by an evolutionary response by the second population to the change in the first ... ‘Diffuse coevolution’ occurs when either or both populations in the above definition are represented by an array of populations that generate a selective pressure as a group (Janzen, 1980).

Janzen’s definition has been extended to include spatiality. Thompson’s “geographic mosaic theory” considers local coevolutionary interactions on a background of geographically distributed species. As such, the geographic mosaic theory allows a diverse set of interactions that are overlooked by Janzen’s definition. In particular, where Janzen specifies pairwise and diffuse interactions, Thompson is keen to acknowledge the multitude of potential interactions that fall between these “catchall” definitions.

[The] geographic view of interactions allows us to replace the catchall phrase ‘diffuse coevolution’ with more-specific, testable hypotheses on reciprocal change among groups of species (Thompson, 1994, p2).

Thus, the generalisation of diffuse coevolution is replaced with the more specific conceptualisation of a set of spatially distributed, locally interacting (and potentially overlapping) pairwise coevolutionary systems. Taken together, Janzen (1980) and Thompson (1994) represent the “standard” definition of coevolution.

Unfortunately, whilst the standard definition effectively restricts the boundary of coevolution, it excludes interactions that many biologists would ideally like to consider (see below). Indeed, Janzen’s strict coevolution is unlike the vast majority of interactions found in nature. As such, his definition is often relaxed in order to encompass a greater variety of interactions. The majority of definitions used by biologists, either implicitly or explicitly, can be covered in this way (Futuyma & Slatkin, 1983, introduction).

In general, artificial coevolution is classified either by vague notions of competitive fitness environments (for example, Hillis, 1990; Angeline & Pollack, 1993; Reynolds, 1994), or alternatively, as the simultaneous evolution of multiple populations with coupled fitness (e.g., Kauffman & Johnsen, 1991; Rosin & Belew, 1997; Nolfi & Floreano, 1998). However, nothing approaching the preciseness of Janzen’s definition exists for artificial coevolution; indeed, when a strict classification appears necessary, Janzen’s definition is occasionally borrowed (e.g., Bullock, 1995). Yet, derived specifically for biological coevolution, it should come as no surprise that Janzen’s definition fails to classify synthetic coevolution sufficiently.

### 3.2.2 The Coevolutionary Continuum

In this section, the standard definition of coevolution is repeatedly relaxed in order to produce a “continuum of coevolution”, with strict coevolution at one end and strict non-coevolution at the other (see figure 3.1). The standard definition with no relaxations is situated at one extreme (far left, figure 3.1); relatively few such systems exist. Each time the definition is relaxed, more systems are included. Once fully relaxed, systems ideally considered non-coevolutionary become encompassed (far right, figure 3.1).

Using the standard definition of coevolution (section 3.2.1), a coevolutionary system must adhere to Janzen’s strict requirements: there must be two populations, with each reciprocally evolving specific adaptations and counter-adaptations in response to the other. Utilising Thompson’s extension, coevolution can also come in the form of local pairwise interactions between geographically distributed populations; with each popula-

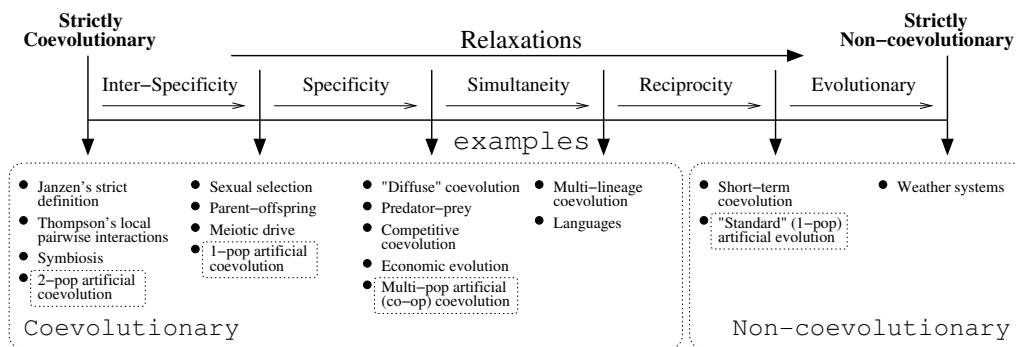


Figure 3.1: The coevolutionary continuum. By repeatedly relaxing Janzen’s (1980) definition, a continuum between strictly coevolutionary and strictly non-coevolutionary systems is formed. Example systems are positioned “as far left” as possible; artificial systems are boxed. Once reciprocity is relaxed, systems are no longer considered coevolutionary (see section 3.2.3). Refer to text for details.

tion potentially engaged in multiple (local) pairwise interactions throughout their spatial range. Systems fitting the standard definition of coevolution are classified “strictly coevolutionary”. On the continuum of coevolution (figure 3.1), no system exists that is “more” coevolutionary than these.

### Strictly Coevolutionary

Given the rigorous requirements of the standard definition of coevolution, very few systems can be classified as “strictly coevolutionary”: only symbiosis and two-population artificial coevolution fit comfortably.

Encompassing parasitism, commensalism and mutualism, symbiosis has been defined as “a condition in which organisms of different species live together in a state of mutual influence” (Ehrman, 1983). The term “symbiosis” implies a period of ongoing interactions (perhaps for an entire lifetime) between two organisms; unlike predator-prey, plant-herbivore and competitive coevolution in which organisms only interact briefly.

[The most intimate of the symbiotic associations] is endosymbiosis; in which one of the organisms incorporates the other—an intricate co-habitation of two species in which a symbiont exists in the cells of a host for at least a discrete portion, if not all, of its partner’s life cycle (Ehrman, 1983).

**Parasitism.** A symbiotic relationship is parasitic if one species lives at the expense of another. Batesian mimicry (Bates, 1862)—where a palatable species mimics the coloration of an unpalatable species in order to gain protection from predation—is a well documented example of parasitism and can be observed in many species of butterfly and snake. Whilst the mimic gains from the protection of coloration, the model suffers from greater predation.

**Commensalism.** If one species benefits from a symbiotic association whilst the other is neither harmed nor benefited, then commensalism has occurred. Having no effect upon one species, commensalism can be difficult to prove. Most interactions have some effect, however small. One example given in nature is the mutual grooming or licking between termitophilous beetles and termites; the beetles having become strongly dependent on the microclimate and food sources provided by the termite nest. This relationship is regarded as commensal rather than mutualistic because termitophiles are relatively rare, with many nests lacking them entirely (Mitter & Brooks, 1983).

**Mutualism.** If both species benefit from association, then symbiosis is mutualistic. Like parasitism, mutualism is much more pervasive than commensalism, and can often result in highly specialised adaptations. A convincing example of mutualism can be found in lichens, which consist of two organisms, an alga (the phycobiont) and a fungus (the mycobiont), in close association.

Lichens can be found in virtually all terrestrial habitats from the polar regions to the tropics [often in] habitats where algae and fungi would find difficulty in surviving alone. (Barrett, 1983).

**Endosymbiosis.** The intimate nature of endosymbiosis can produce such specialised adaptations that species can eventually merge into one lineage. It has been postulated that eukaryotes (organisms with nucleated cells) emerged through the incorporation of free-living prokaryotes (now represented by mitochondria and chloroplasts) by a host cell (Margulis, 1970). In evolutionary terms, this example of endosymbiosis is immensely important (previously mentioned in section 1.3.1).

### **Relaxing Inter-Specificity**

Strict coevolution demands that coevolving traits must lie within different populations, species, or reproductively isolated groups. Relaxing interspecificity allows traits to coevolve within one species, or population, and thus includes coevolution between the sexes (sexual selection), parent-offspring conflict, the sub-organismic coevolution of genes (e.g., meiotic drive), and single-population artificial coevolution through self-play.

**Sexual selection.** Sexual selection, or coevolution of the sexes, occurs between traits that are expressed *either* in females *or* males, but never both. While male and female traits are represented on the same genome, they are never simultaneously expressed: sexual dimorphism results. The classic example of sexual selection

is the peacocks tail. Displaying sexual quality, the peacocks exuberant tail has coevolved in response to female mate preference.

**Parent-offspring conflict.** A more counter-intuitive example of intra-specific coevolution is parent-offspring conflict (see Parker & Macnair, 1979; Parker, 1979). Here, offspring coevolve antagonistically with parents; demanding more investment than parents can afford give. In response, parents counter-act offspring by suppressing their demand. Since “parent” and “offspring” traits are never expressed simultaneously, coevolution can occur.

**Sub-organismic coevolution.** Biological systems (and to some extent, the artificial systems that are based upon them) are inherently hierarchical: genes determine protein structures which in turn form cells; a collection of cells make up an organism; related organisms form a family or kin group and the collection of all kin groups make up a species. Due to the complexity of interactions between each layer, it is virtually impossible to consider one in isolation without losing vital information. Thus, whilst evolution has (at some time) been said to act either at the level of the gene, the individual, kin, social group and species, none of these are *exclusively* correct. In some respect, evolution acts at all levels.

Below the level of the individual, the genetic replicators have been considered the fundamental unit of selection (see Dawkins, 1976, 1982). However, genes are assessed not directly, but via the phenotype of the individual (or “vehicle” in Dawkin’s terminology) that they inhabit. Given the complex phenomena of dominance (one of a pair of heterozygous alleles suppresses the expression of the other), pleiotropy (one allele affects multiple phenotypic traits) and epistasis (loci interact such that one may mask the effect of another), at the level of the genes, cooperation, competition, and antagonism abounds. To be successful, an allele must be able to out-compete alleles challenging for the same loci whilst cooperating with others. Hence, whilst a system may appear evolutionary at the level of the individual, from a genetic perspective the interaction of alleles may be better considered coevolutionary.

Empirical evidence suggests that genetic evolution can occur without reliance on phenotypic testing; hence, purely “selfish” genes that impart no phenotypic advantage can thrive. *Meiotic drive* genes are one such example. During “normal” meiosis a haploid gamete is produced through the *unbiased* segregation of the diploid chromosome; hence, each allele is transmitted with 50% probability. Meiotic drive genes alter this process such that there is excess transmission of one allele over its

alternative form of heterozygote. Thus, meiotic drive genes that favour their own transmission can successfully propagate without necessarily benefiting the phenotype of the organism they inhabit.

In meiotic drive, there is selection for a particular gene, but possessing that gene need not be a causal factor in the survival and reproduction of organisms. Here is genic selection without an organismic benchmark (Sober, 1984).

### **Relaxing Specificity**

Relaxing specificity—enabling one population, or trait, to coevolve with a suite of several others—allows “diffuse” coevolution (a term first used by Gilbert & Ehrlich, 1975). Diffuse coevolution covers the majority of predator-prey, plant-herbivore and competitive interactions observed in nature. In addition, the evolution of economies and multi-population (usually co-operative) artificial coevolution are both encompassed once specificity is relaxed.

**Predator-prey and plant-herbivore.** Perhaps the most familiar coevolutionary interactions are those of predator-prey and plant-herbivores; in which an asymmetric relationship exists such that one species feeds upon another. These interactions are often non-specific—there may be several species involved. Lions of the Serengeti, for example, prey not only upon wildebeest, but also zebra, buffalo, giraffe, etc., and hence coevolve in response to the adaptations of each. However, predator-prey and plant-herbivore interactions are often *more* specific (involve fewer species) than competitive coevolutionary interactions in which entire ecosystems may be involved. Some detailed studies into the specific interactions between predators and prey of two species have been conducted; for example the coevolution of chemical cues (for predator detection) and foraging behaviour (ambush) between nocturnal, rock-dwelling, velvet geckos and elapid snakes. To reduce scent distribution, snakes remain sedentary in rock crevices for days or weeks, waiting to ambush geckos (Downes & Shine, 1998).

**Competitive coevolution.** In the words of Futuyma (1979, p455), two species may compete if both use at least some of the same resources. Consequently, competition in nature is ubiquitous, and the resulting literature is extensive.

The modern literature on competitive coevolution began with Lack’s (1947) work on a genus of primarily ground-feeding finches in the Galápagos islands—Darwin’s

finches—and remains a prime example of character displacement.<sup>2</sup> Lack observed that on islands inhabited by only one species of this genus, bill depth is approximately 10mm. However, if two (or more) species inhabit an island, the smallest bill depth is approximately 8mm and the next smallest 12mm; the 10mm form does not exist. Thus, with bill depth presumed to be related to the size and thickness of the seeds a finch eats, these results suggest that, as a direct consequence of evolutionary competition, the species are utilising different resources (see Roughgarden, 1983a).

In many competitive coevolutionary systems, it is likely that multiple species are interacting.

**Economic evolution.** Like biological systems, economies evolve as organizations learn, adapt and innovate. The global economy can be modelled as a coevolutionary system containing a suite of national economies, each reciprocally evolving in response to adaptations in the others. In the same way, each national economy can also be considered a “diffuse” coevolutionary system, with multiple coevolving organizations; each an economy in its own right. The multiple-hierarchical levels of coevolution observable in the global economy is reminiscent of that observed in biological systems (refer to sub-organismic coevolution, above).

### Relaxing Simultaneity

The simultaneity constraint requires that both populations or traits evolve at the same time; there should be no temporal gaps between evolution and response. By relaxing simultaneity, long periods of stasis can occur. Given that specificity is also relaxed, multi-lineage coevolution—a discrete series of adaptations and counter-adaptations between unrelated species inhabiting a similar niche—becomes coevolution. Arguably, language evolution falls into this category.

**Multi-lineage coevolution.** Multi-lineage coevolution describes a discrete series of evolutionary adaptations and counter-adaptations each occurring in an unrelated species. For example, today’s plant defences may have evolved in response to insects from the Cretaceous period. However, the feeding habits of many of today’s insects may have evolved much more recently (Futuyma & Slatkin, 1983, introduction). Hence, whilst plant defences and insect feeding habits may have “coevolved” as a series of responses, modern and Cretaceous insects may not be of one lineage—modern

---

<sup>2</sup>Competition between species may produce, at places where their ranges overlap, a displacement in the characters of each species from those of the other species. Such a phenomenon—termed “character displacement” (Brown & Wilson, 1956)—is likely to bring an end to the competitive interaction between species, with each diversifying into a separate niche devoid of competition.

insects are not necessarily direct ancestors of the Cretaceous insects that fed on the same plants. As such, the coevolutionary sequence of adaptation and response is not simultaneous; a long period of stasis occurred during the intervening period between the extinction (or shift in feeding behaviour) of Cretaceous insects, and the evolution of modern insect counter-adaptations.

**Language evolution.** Continually mutating, speciating, and eventually becoming extinct, it is clear that languages “evolve”. Languages adapt in response to other languages they come into contact with (historically, via the movement of people). Some words and grammatical forms are adopted, some are lost and others are mutated: potentially, one language may drive the other extinct. However, such “coevolution” is not continuous. Historically, languages compete during the invasion and occupation of one group by another. Once occupation has ended, languages evolve in relative isolation until another invasion occurs. Such coevolution is nonsimultaneous and potentially multi-lineage. With increasing globalisation, however, the dynamics of language evolution have changed dramatically since the industrial revolution.

### **Relaxing Reciprocity**

Reciprocity demands that *both* traits evolve; it is not sufficient for only one trait to evolve in response to the other. For example, reciprocity discounts evolutionary adaptations to “static” abiotic features of the environment (such as altitude, or temperature), since these features do not counter-adapt in response. Relaxing reciprocity allows short-term coevolution in which one trait, or population, has not had time to reciprocally evolve. For example, consider Batesian mimicry. Over a short time-period, whilst the mimic may begin to converge upon the model, it is possible for the model to be unresponsive—selection pressures may not be strong enough to produce an adequate counter-adaptation in the given time-frame. “Standard” one-population artificial evolution—individuals assessed against a non-evolving fitness function— is considered coevolutionary once reciprocity is relaxed.

### **Relaxing Evolutionary**

The coevolutionary continuum is completed once the standard definition of coevolution is fully relaxed, so as to include non-evolutionary systems. Included in the classification “strictly non-coevolutionary” are non-evolutionary complex systems, such as the weather. Whilst weather systems display interesting dynamics and can be modelling using “populations” of local interactions, they do not “evolve” in any biological sense. There is neither competition for resources, nor any analogue of mutation, selection, or lineage. Classifying such systems as “coevolutionary” is clearly absurd.



To successfully categorize coevolution in a meaningful way, the coevolutionary continuum must be appropriately partitioned.

### 3.2.3 When is it *not* Coevolution?

By repeatedly relaxing the standard definition of coevolution, a coevolutionary continuum is created (figure 3.1). With no relaxations, only tightly coupled pairwise interactions, such as symbiosis, are included. However, when fully relaxed, *every* complex system is included, whether evolutionary or not! Clearly, any workable definition of coevolution must fall somewhere between these two extremes. But where?

Considering artificial evolutionary systems, it is important for “standard evolution” to be distinguished from artificial coevolution (single-pop, 2-pop, or multi-pop). This suggests that, in general, reciprocity should not be relaxed. By maintaining reciprocity (and thus also “evolutionary”), non-evolutionary complex systems become non-coevolutionary. Also, coevolution in general becomes non-coevolution over time-scales so short that populations are unable to make evolutionary responses; and thus more closely resemble features of the abiotic environment.

If simultaneity is required, then multi-lineage coevolution becomes non-coevolution. From the perspective of evolutionary computation, the requirement of simultaneity has no effect upon the classification of artificial systems. As such, from the perspective of EC, simultaneity is irrelevant.

If specificity is required, then “diffuse” coevolution—including multi-population artificial coevolution, often used in co-operative coevolutionary systems—becomes non-coevolutionary. In addition, the biological interactions of predators and prey, plants and herbivores, and competitive coevolution in general, each become non-coevolutionary. This is not ideal. Nor are further requirements of specificity and inter-specificity.

Therefore, from the perspective of evolutionary computation, to classify “not coevolution”, the continuum should be partitioned to the left of reciprocity (highlighted on figure 3.1 by the dotted regions). Throughout this thesis, systems to the left of the boundary shall be considered “coevolutionary” and those to the right “non-coevolutionary”. Hence, if an evolving population is not adapting in response to a reciprocally evolving population, then the evolutionary system is *not* coevolutionary.

In section 3.2.5 this informal distinction is formalised via a new definition of coevolution. First, “fitness propensity” and “fitness realization” must be introduced.

### 3.2.4 Fitness Propensity and Fitness Realization

#### Survival of the Fittest: Rejecting the Charge of Circularity

Paradoxically, “fitness” is a central concept in evolutionary theory and yet it is surprisingly vague. Often, a “shorthand” interpretation—equating fitness with reproductive success—is utilised to avoid the challenging formulation of a consistent definition. Whilst this approach is acceptable in many circumstances, confusion prevails. Consequently, debate over the tautological nature of the phrase “survival of the fittest”—those best at surviving are the ones who survive—has raged for well over a century.<sup>3</sup>

Attacking the “traditional” definition of fitness, Mills and Beatty (1994) reject the charge of circularity by suggesting that the simplified use of “fitness” as a *post hoc* measure of reproductive success is the source of confusion. Clearly, any such definition is purely descriptive; not explanatory. In response, Mills and Beatty (1994) equate “fitness” with the *ability* of an individual to survive and reproduce in a given environment—a propensity analysis of fitness that is non-circular.

We believe that the confusion involves a misidentification of the *post facto* survival and reproductive success of an organism with the *ability* of an organism to survive and reproduce. We believe that “fitness” refers to the ability . . . Thus, we suggest that fitness be regarded as a complex *dispositional* property of organisms. Roughly speaking, the fitness of an organism is its *propensity* to survive and reproduce in a particularly specified environment and population . . . In general, we want to rule out the occurrence of any environmental conditions which separate successful from unsuccessful reproducers without regard to physical differences between them (Mills & Beatty, 1994).

Consider identical twins, *A* and *B*. Imagine both live in the same local environment and interact with the same conspecifics. Each leads a similar existence until one stormy day a chance lightning strike kills (a still immature) twin *A*. Twin *B*, however, is unharmed in the storm and goes on to lead a successful life, eventually siring 10 grandchildren. Under the “traditional” definition of fitness, twin *B* is considered far fitter than twin *A*. However, any “workable” definition must surely equate them: if genetic clones vary in fitness, then what exactly does it tell us? The approach of Mills and Beatty (1994)

---

<sup>3</sup>First introduced by Herbert Spencer (1864, p444-445), the term “survival of the fittest” is superficially circular when fitness is defined purely in terms of survivability and reproductive success. To attribute type *A* as *fitter* than type *B* is purely to say that type *A* on average leaves more offspring than type *B*. However, this cannot be *explained* by the difference in fitness between *A* and *B*, since fitness is defined in terms of reproductive success; i.e., *after* the event (see Mills & Beatty, 1994).

circumvents this problem. As identical twins,  $A$  and  $B$  share an equal fitness ability, or “propensity” to survive and reproduce in their particular environment. The chance death of  $A$  in immaturity does not change this. Irrespective of the eventual outcome, both twins (if placed in the same environment and population) begin with equal reproductive potential.

Thus, using the “dispositional fitness” of Mills and Beatty (1994), “survival of the fittest” can be more accurately rephrased as “genes with the greatest propensity to survive and reproduce in the current environment are those most likely to propagate”. The charge of circularity is rejected.

### Propensity and Realization

Continuing the reasoning of Mills and Beatty (1994), let us dissect fitness into two parts: “fitness propensity” ( $f_P$ ), the complex dispositional property of an organism to survive and reproduce in a given environment; and “fitness realization” ( $f_R$ ), the actual reproductive success of an organism, calculated *post facto*.<sup>4</sup> Then,  $f_P(x)$  is the *expected* reproductive success of  $x$ , determined *prior to* the event, and  $f_R(x)$  is the *actual* reproductive success of  $x$ , determined *after* the event.

With two measures of fitness, it becomes possible to account for contingency. Consider identical twins  $A$  and  $B$  once again. If placed in the same environment, each has identical fitness propensity  $f_P$ . Any variation in fitness realization  $f_R$ —the actual reproductive success of each twin—must be a result of contingent events such as lightning strikes.

In artificial evolutionary systems,  $f_P(x)$  directly equates to the “fitness score” attributed to  $x$  following assessment. Each generation, the probability of reproduction is a monotonically increasing function of fitness (refer to 2.2.1). Hence, the “fitness score” of  $x$  is the propensity to reproduce,  $f_P(x)$ . In contrast,  $f_R$  is calculated by the actual number of offspring; determined by the (deliberately) stochastic process of selection. Artificial “twins”  $A$  and  $B$ , if placed in the same environment, have identical fitness propensity,  $f_P(A) = f_P(B)$ . The stochastic nature of selection, however, makes it possible (indeed probable) that  $f_R(A) \neq f_R(B)$ . Artificial “lightning strikes” and other contingent phenomena “occur” in the selection phase of artificial evolution. Since  $f_P$  is directly observable in artificial systems, using the values of  $f_P(x)$  and  $f_R(x)$ , it is possible to ascertain whether individual  $x$  underachieved or not. In biological systems, this is a much harder proposition.

---

<sup>4</sup>Given Hamilton’s (1964) concept of inclusive fitness, it is perhaps more correct to talk in terms of genes rather than individuals. Refer to section 1.3.5.2.

Notation	Description	Evolution	Coevolution
$f_P(x)$	Fitness “propensity” of individual $x$ in environment $E$	$E$ fixed	$E$ reciprocally evolves
$f_R(x)$	Fitness “realization” of individual $x$ in population $P$	$f_R(x)$ varies with make up of $P$	

Figure 3.2: Fitness propensity and fitness realization in artificial systems. In *both* evolutionary and coevolutionary systems,  $f_R(x)$  varies depending upon the make up of the population  $x$  is in. However,  $f_P(x)$  only varies under coevolution. In standard evolutionary systems the “effective environment” of  $x$  is fixed.

### Evolution and Coevolution

In artificial systems, “fitness propensity” and “fitness realization” can be used to distinguish evolution from coevolution (refer to figure 3.2).  $f_P(x)$  denotes the propensity of individual  $x$  to survive and reproduce in a particular “effective environment”  $E$ . In artificial evolutionary systems, environment  $E$  is determined by the assessment function—the process used to calculate the “fitness score” of individuals. Let  $f_P^{Tr}(x)$  represent the (perhaps unknowable) “true” fitness propensity value of  $x$ , given  $E$ .

In “standard” evolutionary systems,  $E$  is fixed—the assessment of individuals does not vary through time. Thus, for deterministic assessment,  $f_P(x) = f_P^{Tr}(x)$ . However, if assessment is probabilistic,  $f_P(x)$  is distributed with mean  $f_P^{Tr}(x)$ . In artificial evolution, variation in  $f_P(x)$  is a direct result of stochasticity in the assessment function  $E$ . Since  $E$  does not vary,  $f_P^{Tr}(x)$ —the true fitness propensity of individual  $x$  in effective environment  $E$ —is fixed for all time.

Contrastingly, in coevolutionary systems,  $E$  *does* vary—the assessment function reciprocally adapts in response to adaptations in  $P$ , the evolving population. As  $E$  varies, so too does  $f_P^{Tr}(x)$ . Thus, if  $f_P^{Tr}(x)$  varies with respect to a reciprocally evolving  $E$ , one can classify the system as coevolutionary. For *each*  $E$ , stochasticity in assessment ensures that  $f_P(x)$  is distributed with mean  $f_P^{Tr}(x)$ .

The selection process biases reproduction in favour of individuals with (relatively) high fitness propensity  $f_P$ .  $f_R(x)$  denotes the fitness realization (number of offspring) of individual  $x$  in population  $P$ . Due to the stochasticity of selection,  $f_R(x)$  is normally distributed with mean value  $\overline{f_R}(x)$ . In addition,  $\overline{f_R}(x)$  varies with the make up of  $P$ —the population  $x$  is a member of. Depending upon whether  $f_P(x)$ , given  $E$ , is high or low relative to other individuals in the population,  $\overline{f_R}(x)$  will vary accordingly. This is true in *both* evolutionary *and* coevolutionary systems, and is a direct consequence of the struggle

for scarce resources underpinning all evolution.

### 3.2.5 A New Definition of Coevolution

This thesis advances the following “reciprocal” definition of coevolution:

A system is considered coevolutionary *if and only if*  $f_P^{Tr}(x)$ —the “true” fitness propensity of each evolving individual (or trait),  $x$ —varies with respect to other *reciprocally evolving* individuals (or traits).

Considering the continuum of coevolution (section 3.2.2; figure 3.1) the reciprocal definition of coevolution classifies everything to the left of “reciprocity” as “coevolutionary”, and everything to the right as “non-coevolutionary”. Hence, the reciprocal definition is equivalent to the standard definition of coevolution (section 3.2.1) with inter-specificity, specificity and simultaneity relaxed. Reciprocity, however, is still required.

Considering biological systems, the reciprocal definition classifies (amongst others) symbiosis, sexual selection, “diffuse” predator-prey interactions and multi-lineage coevolution as “coevolutionary”. Of artificial systems, single-, 2-, and multi-population competitive (and cooperative) coevolutionary systems are all classified “coevolutionary”. “Standard” evolutionary systems, however, are not. The reason for this can be highlighted by considering an evolving system containing only one population,  $P$ .

If individuals within  $P$  are assessed via a static function, then  $f_P^{Tr}(x)$  of each individual,  $x$ , is independent of every other individual. Whilst  $f_R^{Tr}(x)$  is free to vary depending upon the make up of  $P$ ,  $f_P^{Tr}$  is not. As such, the system is “non-coevolutionary”. If individuals within  $P$  are assessed through competition with either fixed, or randomly chosen (non-evolving) opponents, then the system is again classified “non-coevolutionary”. Whilst  $f_P(x)$  may vary with opponents, opponents *do not* evolve in response. Hence,  $f_P^{Tr}(x)$  does not vary—the effective environment,  $E$ , does not evolve. However, if individuals within  $P$  are assessed through competition with other (evolving) members of  $P$ , then the system *is* “coevolutionary”— $f_P^{Tr}(x)$  varies with reciprocally evolving opponents.

### 3.2.6 Relativism versus Absolutism

Consider an artificially evolving population,  $P$ . Then  $\forall x \in P$ , if  $f_P^{Tr}(x)$  varies relative to other (reciprocally) evolving individuals, let us say that the fitness propensity of  $x$  is “relative”; denoted  $f_P^{rel}(x)$ . Else—if the effective environment of  $P$  is not reciprocally evolving—say that the fitness propensity of  $x$  is “absolute”; denoted  $f_P^{abs}(x)$ . Using the reciprocal definition of coevolution (page 3.2.5), one can now classify systems that assess

relative fitness propensity ( $f_P^{rel}$ ) as coevolutionary, and those assessing absolute fitness propensity ( $f_P^{abs}$ ) as non-coevolutionary. This terminology is similar to that generally used in the EC literature, since “relative fitness” and “absolute fitness” are sometimes used to distinguish coevolution from non-coevolution (e.g., Watson & Pollack, 2001). However, such terms are vague (and potentially ambiguous) and thus benefit from the rigorous treatment of fitness performed throughout section 3.2. Henceforth, the concise terminology of “relative fitness” (meaning “relative fitness propensity”,  $f_P^{rel}$ ) and “absolute fitness” (meaning “absolute fitness propensity”,  $f_P^{abs}$ ) shall be used throughout this thesis.

Absolute fitness functions measure fitness against some objective *yardstick* that does not vary over time. However, this approach becomes difficult when the qualities that are desired in an optimal solution are poorly understood, difficult to define ahead of time, or hard to operationally assess.

Faced with these problems, it is often attractive to define fitness in relative terms. Rather than create some yardstick metric, it is sufficient to define rules of contest that award individuals the status of winner and loser. Search should then be biased in favour of winners and against losers. Rather than attempting to define and measure the quality of being a good tennis player, one merely needs to count the number of victories that a player achieves, and use this as a proxy for tennis-playing ability.

It is quite possible to prepare, ahead of time, a fixed set of opponents. This has the advantage of allowing comparison between different players on the basis of how they fare against this test-set. However, for this approach to work, it is critical to choose effective opponents against which a player is assessed. A fixed test-set must be capable of effectively discriminating all the players that the search process encounters. Two kinds of problem immediately arise. First, if the test set is too easy, or too difficult, it may not discriminate between players of differing quality. Second, if the test-set exhibits a bias (the players are all right-handed, for example), tested players may achieve many wins through exploiting this bias, rather than approximating an optimal generalist strategy. It turns out that determining the properties of an effective fixed test set in advance is equivalent to the problem of defining an absolute fitness function.

Considering this, the term “relative fitness measure” is reserved for situations in which the test set evolves in response to the performance and character of the players we are searching through. For example, letting one evolving population provide the test set for a second, and *vice versa*, ensures that measures of success are truly relative.<sup>5</sup> Each in-

---

<sup>5</sup>This approach was used by Hillis (1990, refer to section 3.1). However, for the definition of relative fitness used throughout this thesis, it is acceptable for there to be only one evolving population from which

dividual's quality is judged on the basis of its performance relative to a set of transient opponents—a group of contemporaries that may never be reassembled during the rest of the search process.

This set up cleverly avoids the second of the problems mentioned above. Any exploitable bias that might have stymied a standard evolutionary search algorithm is unlikely to remain around for very long. Individuals who are exploited in this way lose contests and are selected out of the population as a result. However, the coevolutionary approach does nothing to prevent the first problem—that of playing a set of opponents that are either all too easy or all too difficult—and, moreover, brings with it several new problems (see section 3.4).

As a warning to the reader, the terms “absolute” and “relative” used in this section should not be confused with the mathematical modelling usage often employed in classic models of evolution; described below.

Let  $l$  be the probability that a zygote survives to adulthood and let  $m$  be the fertility of an adult (where  $m$  equals half the number of gametes that eventually become incorporated into zygotes). Then the *absolute* fitness, in numbers of individuals, is  $W = ml$ , where  $w \geq 0$ . *Relative* fitness is calculated as  $w = \frac{ml}{W_{max}}$ , where  $W_{max}$  is the highest absolute fitness that any individual has, hence  $1 \geq w \geq 0$  (see Roughgarden, 1983b).

### 3.2.7 Moving Mountains: A Dynamic Landscape

Since its introduction in the 1930s the concept of a fitness landscape has been a useful tool for visualizing evolutionary processes (Wright, 1932, 1967). Assuming that every point on a hyper-plane denotes a possible genotype, then attributing a height proportional to the fitness propensity of the genotype at each of these points produces a fitness landscape with good individuals residing near the tops of peaks and poor individuals in valleys. An evolving population can then be thought of as something like a set of competing hill-climbers.

Under an evolutionary framework, the fitness landscape is considered static, since the mapping from genotype space to fitness scores is constant. It makes sense to represent a genome's quality as a point on the fitness landscape, since this landscape depicts the nature of certain constraints that are imposed on the outcome of competition between genomes. This is despite the fact that the reproductive success of a genome is dependent not solely on its fitness, but on the relationship between this value and the fitness scores

---

opponents are drawn.

of those genomes with which it is currently in direct competition.

Consider two genomes,  $g_0$  and  $g_1$  with fitness scores  $f_P^{abs}(g_0)$  and  $f_P^{abs}(g_1)$ , respectively. At different times, these same fitness scores may translate into different degrees of reproductive success— $f_R(g_0)$  and  $f_R(g_1)$ —depending on the make-up of the population that  $g_0$  and  $g_1$  find themselves in. However, under a standard evolutionary scheme, if  $f_P^{abs}(g_0) > f_P^{abs}(g_1)$ , then  $g_0$  will tend to enjoy greater success than  $g_1$  *whatever* the make-up of the population. The shape of the fitness landscape imposes an *ordering* on the success of genomes it depicts, without specifying the magnitude of this success. The altitude of a single point in the landscape provides no information concerning the reproductive success that might be experienced by the associated genome. However, a *difference* in altitude between two points on the landscape represents the relative difference in reproductive success that these genomes would experience should they ever directly compete.

Under a coevolutionary regime, this interpretation must be radically reworked. Since relative fitness measures ensure that the mapping  $G \mapsto F$  from genotype space to fitness values is free to vary over time, it is quite possible that  $f_P^{rel}(g_0) > f_P^{rel}(g_1)$  at time,  $t_0$ , whilst at some other time,  $t \neq t_0$ , the opposite inequality holds. As a result, no single static fitness landscape can depict the distribution of fitness (propensity) values across the space of possible genotypes for all time. Rather, the landscape must be represented by a dynamic scene, continually undulating as fitness values change. Thus, continually ebbing and flowing, the coevolutionary fitness “landscape” is perhaps better analogised as a fitness “seascape”.

The ability of an evolving population to reach optimal or near-optimal solutions is dramatically affected by the shape of a fitness landscape. The ideal landscape for evolution contains one high peak with gently sloping sides, such that from any initial position the evolving population is able to climb a continuous gradient towards the desired solution. This is sometimes termed a “Mt. Fuji” landscape. In contrast, the “worst” landscapes for evolution are highly rugged with many small peaks. With little correlation between mutationally adjacent genotypes, the evolving population wanders aimlessly, continually changing direction (and repeatedly getting stuck). In general, landscapes fall somewhere between these two extremes.

When setting up an evolutionary algorithm, the static fitness landscape is entirely determined by the particular assessment function, genetic encoding and genetic operators chosen. Unfortunately, understanding *how* to mould the desired fitness landscape for a particular problem is often as difficult as solving the problem itself. Hence, guesswork and chance play a major role in the success of any evolutionary algorithm.



In comparison, whilst implementation choices also play a key role in coevolutionary algorithms, the continual deformation of the landscape during run-time dramatically reduces the significance of such design issues. In conjunction, dynamic landscapes offer a means of circumventing the traditional problem of premature convergence at a local optimum. Once an evolving population reaches a local peak in the static landscape, it can only escape by a large mutation that can effectively jump the population across a fitness valley and onto an alternative (higher) peak.<sup>6</sup> Obviously, the greater the mutational distance, the more unlikely the required mutation is. Although practical methods for tackling premature convergence abound (see Goldberg, 1989), the problem still receives much attention from the EC community.

The dynamic landscapes of coevolution circumvent the problem of local optima; with fitness peaks being continually eroded, local optima often exist for only short periods. Thus, the dynamic landscape of coevolution offers a potentially great advantage over evolution. However, whilst there is reason to believe that the move to coevolution is associated with some attractive benefits, it also introduces a number of important problems that threaten to overwhelm them. These are detailed fully in section 3.4.

### 3.2.8 The Nature of Objective Fitness

Depending upon the presence, or absence, of some (external) objective measure of fitness, *all* coevolutionary systems fall into one of the following two categories:

- 1. Although it may be difficult to define, *some* objective fitness measure exists.** When engaged in coevolutionary optimisation, computer scientists *always* have some concept of objective fitness. The very nature of optimisation means that it must be true that some “solutions” are regarded as “better” than others.
- 2. *No* objective fitness measure exists.** In nature, there is *no* concept of objective fitness: the fitness of individuals is implicitly realised through life-time survival. Thus, one cannot assert value-judgements on the outcome of biological evolution: if a species becomes extinct, for instance, then *objectively* it is neither a “good” nor “bad” outcome.

---

<sup>6</sup>An alternative paradigm of “neutral networks” has been proposed; rather than hill-climb across the landscape, a population is thought to drift along connected networks of equal fitness. Such networks permeate genetic space. Through occasional jumps between networks it is thought that, given enough time, any possible fitness value can be obtained. If correct, the scenario of a population trapped on a local peak disappears (Barnett, 1998, also refer to chapter 8).

To some readers it may appear that, between these two “extremes”, there is a third category covering the hazy middle-ground—perhaps including some open-ended artificial life simulations. However, it is the author’s belief that no middle-ground exists; only confusion. This confusion arises because researchers haven’t decided which category their simulations fall into. They may hope to just run the simulation and see what happens, however, if there is no concept of objective fitness, then it becomes impossible to talk about “disappointment” when regarding results. Before designing a simulation, practitioners need to decide upon which side of the fence they will sit.

### 3.3 The Success Stories

Having first been applied to problem solving in the early 1990s, coevolutionary algorithms have rapidly grown in popularity. In this section a brief (and by no means exhaustive) review of the more successful coevolutionary applications is given. Hopefully, this should give the reader an appreciation of coevolution’s current versatility and pervasiveness.

Hillis’ (1990) seminal work proved by example how artificial coevolution can be applied to optimisation problems. Hillis chose to design minimum comparison sorting networks, a domain with a long and competitive history that has resulted in the current record of sixty comparisons to sort sixteen elements, achieved *by hand* in 1960 by M. W. Green (Knuth, 1973). Initially, using randomly generated inputs as test sets, Hillis evolved the sorting networks on a static fitness landscape, resulting in a minimal network of 64 comparisons. He found that two factors were preventing the evolution of shorter networks. Firstly, the classical problem of local maxima made it difficult for the system to progress once a reasonable solution had been found. Secondly, the test process proved inefficient—after a few generations most of the inputs were fully sorted by the vast majority of networks. To compensate for these problems, Hillis allowed the test cases to coevolve with the sorting networks, giving them complementary fitness functions and thus producing an artificial host-parasite system. This allowed the lists to evolve towards punishing the weaknesses of suboptimal networks, whilst dramatically reducing the number of redundant tests—those *too easy* to sort. Using this method, Hillis discovered a 61-comparison network, a much better solution than previously evolved and only one comparison longer than the best known solution.

Hillis’ success has since inspired a whole spectrum of research utilising coevolution for practical problem solving. Though coevolution has been applied to a diverse range of domains, sorting networks have continued to attract attention (e.g., Juillé, 1995; Olsson, 1996; Rosin, 1997). Most impressively, Juillé’s Evolving Non-Determinism (END)

model improved upon a 25-year-old record by discovering a network of 45 comparisons for list inputs of 13 elements. Currently, this remains the shortest known solution (Knuth, 1998).

In 1994, Sims (1994a, 1994b) coevolved a population of virtual creatures in a physically simulated three-dimensional environment by engaging individuals in pairwise competition for a common resource. In this pioneering work, both morphology and neural controllers (encoded as directed graphs) were genetically determined, allowing bodies and behaviour to coevolve simultaneously. As a result, Sims produced a set of oddly familiar morphologies and behaviours that included walking, swimming, hopping and following. Further success on body-brain coevolution has been achieved by Hornby and Pollack (2001, 2002). Rather than encapsulate the genetic encoding as directed graphs, Hornby and Pollack utilise the generative encoding of L-systems (originally introduced by Lindenmayer, 1968, to model the development of relatively simple multicellular organisms, such as algae), producing hierarchies of regularity which allow reuse. The resulting designs have been transferred successfully into reality (Hornby, Lipson, & Pollack, 2001).

Coevolution has been most widely applied to game-playing. The Iterated Prisoner's Dilemma (IPD) is a two-player non-zero-sum game designed to study cooperation in game theory and extensively applied in politics, economics, psychology, biology and AI (Axelrod, 1984). Each round, players have a choice of either defecting or cooperating; whilst the greatest total payoff comes from mutual cooperation, the temptation each round is to defect. Coevolutionary algorithms have been successfully applied to the IPD, reaching the unstable equilibrium of mutual cooperation (Axelrod, 1987; Lindgren & Nordahl, 1994). Coevolution has also been applied to more realistic extensions of the IPD: the continuous IPD, where a range of intermediate choices between full co-operation and full defection are offered (Darwen & Yao, 2000, 2001); and the multiple-player NIPD (Yao, 1997). Following TD-Gammon (Tesauro, 1992)—a backgammon program that, through self-play, became one of the best players in the world—backgammon strategies have been successfully discovered using simple coevolutionary hill-climbing algorithms (Pollack et al., 1996; Pollack & Blair, 1998). However, it is thought that the fantastic success of coevolution in this domain may be due to the inherent properties of backgammon, especially the stochastic dice roll and *reversibility* of the game. As such, whilst coevolution has achieved fair success when applied to other games, the results have not been as outstanding as those with backgammon. Those worth mentioning include the interactive video-game Tron (Funes, Sklar, Juillé, & Pollack, 1998; Funes, 2001), the ancient Chinese board game Go (Rosin, 1997; Lubberts & Miikkulainen, 2001), checkers (Fogel, 2000), Tic-Tac-Toe (Angeline & Pollack, 1993; Rosin & Belew, 1997), Texas Hold'em

Poker (Noble & Watson, 2001) and Nim (Rosin & Belew, 1997).

The pursuit and evasion of predators and prey has been a successful (and popular) domain for coevolutionary research; in abstract (Koza, 1991; Ficici & Pollack, 1998b), in agent simulation (Miller & Cliff, 1994b; Reynolds, 1994; Cliff & Miller, 1995, 1996) and in physical robots (Floreano & Nolfi, 1997a, 1997b; Floreano, Nolfi, & Mondada, 1998; Nolfi & Floreano, 1998). Coevolution has been successfully applied to classification tasks, including the neural network pattern classification of intertwined spirals (Juille & Pollack, 1996) and the density classification problem of 1-D cellular automata (Paredis, 1997; Juillé & Pollack, 1998a, 1998b; Ficici & Pollack, 2001), as well as function optimisation (Potter & De Jong, 1994; Paredis, 1995; Pagie & Hogeweg, 1998). Finally, other interesting fields include the development of communication (Werner & Dyer, 1992; Ficici & Pollack, 1998b), job shop scheduling (Husbands & Mill, 1991), room painting (Puppala et al., 1998) and robot football (Østergaard & Lund, 2003).

## 3.4 Critique

By utilising relative rather than absolute fitness assessment, coevolution manages to circumvent some of evolution's more debilitating problems (refer to section 3.2.6). However, the extra dimension of dynamical complexity introduced through relative fitness is a double-edged sword. While there is reason to believe that this move is associated with some attractive benefits, it introduces a number of important problems that threaten to overwhelm them. These are detailed below.

### 3.4.1 Visualization

Though coevolutionary systems exhibit several unwanted behaviours (see sections 3.4.2-3.4.5), the predominant problem regarding coevolution does not directly involve a system's dynamics. Rather, when utilising coevolution for problem solving, perhaps the greatest challenge comes when attempting to accurately visualize a system's dynamical behaviour. It is inherently difficult to visualize progress in a coevolutionary system because of the relative fitness functions they incorporate (section 3.2.6); the move from absolute to relative fitness immediately renders some standard visualization techniques redundant. However, deciphering the trajectory of coevolution *blindly* (i.e., without an adequate visual tool) is highly problematic. Thus, in the implementation and analysis of coevolutionary systems, visualizations are a key component; in order to avert the onset of problematic dynamics, it is necessary to firstly identify their occurrence.

The progress of an evolutionary algorithm has traditionally been measured using the graphical representation of fitness over time. This appears legitimate when considering evolution: if average population fitness increases over time, we can be reasonably confident that random individuals chosen from later generations will tend to be of higher quality than their ancestors. However, when considering coevolutionary algorithms, graphs of fitness over time are far more problematic. It is widely recognised that the manner in which the distribution of fitness changes over generational time is no indicator of particular trends in population quality. Increases in fitness may be the result of one population improving at the expense of the other, or both populations improving (or regressing) at different rates. Indeed there is no guarantee that a trend in quality will be reflected in any way by a population's fitness scores. This stems from the fact that fitness values recorded at generation  $t$  are incomparable to those recorded at some other time; the variable plotted for generation 0 is simply not the same variable as the one plotted for generation 1. To label the ordinate of such a graph "fitness" is to invoke a non-category; unsurprisingly, such graphs are uninformative and alternatives have been sought.

The simplest way to visualize coevolutionary progress is to plot some externally imposed objective measure over time. This method was used by Hillis (1990) during his seminal work coevolving sorting networks; at the end of each generation, the elite sorting network (that with the greatest sorting success and, hence, highest relative fitness) was exhaustively tested by the complete set of potential input lists. Progress was then monitored by plotting the percentage of all lists the elite network from each generation was able to sort. However, whilst this approach worked well for the domain of sorting networks, such an evaluation function demands a finite, discrete space of test inputs that can be exhaustively generated in a reasonably short period of time. Thus, in complex, real-world domains, the option of exhaustive testing is often unavailable.

As an alternative, one may contrive an objective measure by competing, or testing, coevolved solutions against a (small) *fixed set* of benchmark trials. This approach was successfully used by Pollack et al. (1996) to test the ability of their coevolved backgammon strategies; Tesauro's expert linear evaluation function "PUBEVAL" was used as a yardstick. By removing the strict requirements of exhaustive testing, using fixed opponents is a much more versatile alternative. However, this method will only measure progress as long as the test set can adequately discriminate between the evolved solutions. In conjunction, it is possible that the design of a suitable test-set may be as difficult as the problem one is trying to coevolve solutions for; this is particularly true if the system is symmetrical. As such, this method will only suffice in a small subset of domains.

In the absence of an objective yardstick, entirely new visualizations become neces-

sary. Often, these are created *ad hoc* and rely heavily upon the idiosyncrasies of the specific domain they were designed for. As such, very few generic visualization tools exist. Evolutionary activity plots (originally Bedau & Packard, 1992) are a possible exception. Although not originally designed for coevolution, these plots are able to visualize the rate at which useful genetic innovations are absorbed into an evolving population without requiring an objective assessment function. By tracing the lineage of adaptively significant genotypes, evolutionary activity plots highlight useful genetic adaptations from periods of random drift. However, whilst evolutionary activity plots afford insight into the underlying genetics of a coevolutionary system, they fail to sufficiently highlight coevolutionary progress in the necessary manner.

As far as the author is aware, the first generic tool specifically designed to visualize progress in coevolutionary systems was the “current individual against ancestral opponent” (CIAO) plot (Cliff & Miller, 1995). Entirely removing the need for an objective measure, CIAO plots display the relative performance of the elite individuals from each generation; as coevolution unfolds, the best individual from each generation is pairwise assessed against the best individual from each ancestral generation. The results of each competition are plotted on a 2-dimensional grey-scale matrix, where darkness represents success and the co-ordinate of each point reflects the generation of each elite competitor (see section 7.2.1). Hence, CIAO plots visualize the *relative* quality of coevolved solutions, and are thus able to distinguish coevolutionary progress (smooth gradient from light to dark) from cycling (diagonal banding) and stasis (a homogeneous plot). Since their introduction, CIAO plots have been successfully used in a number of domains; however, they suffer from unexpected ambiguities and should be used in conjunction with other techniques (refer to chapter 7).

In recent years “dominance plots” have been proposed as an extension of the CIAO plot (Stanley & Miikkulainen, 2002). Using information from ancestral-opponent contests, dominance plots are drawn by marking generations in which the current elite *dominates* all previous dominant elites.<sup>7</sup> The resultant graph produces a “bar-code” style pattern, with each black line representing the appearance of a new dominant individual.

Let  $d_j$  be the  $j$ th dominant strategy to appear over evolution. Then dominance is defined recursively as follows:

- $d_1$  is the elite of generation 1
- where  $j > i$ ,  $d_j$  is an elite such that  $\forall i < j$ ,  $d_j$  beats  $d_i$

---

<sup>7</sup>Here, “dominance” is similar, but not identical, to pareto dominance; section 3.4.2.

Dominance plots are used to identify a sequence of increasingly more sophisticated strategies. By ignoring comparisons between non-dominant elites, noisy fluctuations in progress are ignored. Whilst this offers an advantage over CIAO plots, dominance plots discard so much information that they risk throwing the baby out with the bathwater. For caution, practitioners are advised to use dominance plots in tandem with CIAO plots, rather than as an exclusive alternative.

### 3.4.2 Overspecialisation

Throughout the evolutionary process, adaptive traits are gradually acquired over time, each a specific adaptation to the environment they inhabit; as a result, populations generally become highly specialised. However, whilst this phenomenon has produced the breathtaking abundance of biological diversity we observe around us, when utilising co-evolution (or evolution) for problem solving, specialisation can prove problematic.

The subjective concept of *overspecialisation* occurs when coevolving populations become “too focused”. By exploiting the idiosyncratic weaknesses of their opponents, populations may evolve in an unexpected (and perhaps unwanted) direction. For example, when coevolving tennis players, one population may begin to exploit the slow reactions of the other population by hitting the ball harder. Through time, there may be a continuous escalation in the power of individuals at the expense of all other facets of the game, such as agility, accuracy and stamina; the resulting tennis players are left ill-equipped to tackle more generalist *all-round* strategies. Overspecialisation potentially results in brittle solutions that are unable to generalise (Darwen & Yao, 1995; Watson & Pollack, 2001; Bucci & Pollack, 2002, 2003).

Whilst coevolutionary interactions vary greatly in specificity, from fairly generalist predator-prey interactions to highly specific endosymbiosis, it is generally erroneous to characterize a species as *overspecialised*. Unless specialisation directly results in catastrophe (extinction), one cannot regard it as problematic; the inherently value-free nature of biology makes it impossible for us to determine what *is* and *is not* overspecialised (refer to 3.2.4). As such, the problem of specialisation exists purely within artificial coevolution and has no discernable analogue in biology.

The problem of overspecialisation is associated with the more traditional machine learning problem of over-fitting.<sup>8</sup> In coevolution, over-fitting (overspecialisation) may be avoided by maintaining diverse training and rule sets (the evolving populations) and is

---

<sup>8</sup>A machine learning algorithm is said to over-fit when it adapts so well to a *specific* training set that the learning model treats noisy idiosyncrasies in the data as meaningful.

achieved by maintaining genetic diversity within the coevolving populations (Darwen & Yao, 1997).

Evolutionary computing is littered with diversity maintenance (or niching) techniques (for a review, see Goldberg, 1989). Approaches range from standard and deterministic crowding (where competition is restricted to genetically similar individuals; De Jong, 1975; Mahfoud, 1995, respectively), to spatial embedding (where competition is restricted by spatial locality; first used in a coevolutionary context by Hillis, 1990) and islands models (where populations are split into multiple demes that interact through migration only; originally Wright, 1932). Fitness sharing (first proposed by Goldberg & Richardson, 1987) has been perhaps the most successful; fitness is shared among genetically similar individuals, thus rewarding those that are genetically unique. However, fitness sharing is sensitive to the “radius of similarity”; a parameter defining the volume of genetic subspace deemed similar. As such, fitness sharing has been extended to remove the definition of a radius parameter (Smith, Forrest, & Perelson, 1993). Here, individuals “bid” for the pay-off of discrete objects; those scoring the highest sharing the reward. Further extensions have enabled fitness sharing to be applied to coevolution: an individual receives a payoff when it has the best score against an opponent; draws are shared. First introduced as “competitive fitness sharing” (Rosin & Belew, 1995),<sup>9</sup> this technique has also been adopted as “implicit sharing” (Darwen & Yao, 1996) and “resource sharing” (Juillé & Pollack, 1998b); however, the details of each are similar. Probably the most descriptive term, *resource sharing* will be used throughout this thesis.

Resource sharing maintains genetic diversity in a population by encouraging niching; individuals are rewarded for being able to solve tests that few others can. This idea has been extended to coevolutionary scenarios where opponents are treated as a commodity or resource. Rather than gain a fitness point for each victory against an opponent (simple fitness), one fitness point is shared among all the competitors that beat a particular individual. Thus, individuals are rewarded less for *how many* opponents they beat and more for *who* they beat; rewarding phenotypic diversity and maintaining genetic diversity. By niching the population into specific sub-populations, resource sharing not only counteracts focusing, but also raises the potential for multi-objective optimisation (MOO); whilst each niche is inherently specialised, the population as a whole forms a generalist set capable of mastering several objectives.

---

<sup>9</sup>In conjunction with competitive fitness sharing, Rosin and Belew (1995; Rosin, 1997) also proposed *shared sampling*; a method for choosing diverse opponents *worth beating*. Rather than sample opponents at random, shared sampling purposely selects a sub-set of parasites that challenge all segments of the host population; based upon the previous generation’s performance. Shared sampling may produce some of the same effects as competitive fitness sharing.



As a recently introduced alternative to niching, pareto coevolution<sup>10</sup> offers a way of preserving useful (though perhaps idiosyncratic) adaptations by explicitly treating the performance against each player in the population as a dimension for MOO (Ficici & Pollack, 2001; Noble & Watson, 2001; Bucci & Pollack, 2003; De Jong & Pollack, 2004). Individuals compete against all others in a set of pairwise competitions. Any individual that is pareto-dominated by another (i.e., scores equal or less against all opponents) is destroyed; those that are non-dominated are preserved. Novel individuals are created from parents randomly chosen from the front.

**Definition (Pareto-Dominance):** Let  $\mathbf{a} = (a_1, \dots, a_k)$  denote the scores of individual  $A$  against each of  $k$  opponents (objective dimensions) and let  $\mathbf{b} = (b_1, \dots, b_k)$  denote the scores of individual  $B$  against the same  $k$  opponents (objective dimensions); then  $A$  pareto-dominates  $B$  if and only if  $A$  performs at least as well as  $B$  in all dimensions and better than  $B$  in at least one, hence:

$$\mathbf{a} \succ \mathbf{b} \iff \forall_i \in \{1, \dots, k\}, a_i \geq b_i \wedge \exists_i \in \{1, \dots, k\} : a_i > b_i$$

By preserving all non-dominated individuals, pareto coevolution ensures that any potentially useful adaptation is not lost; making it unlikely for the population as a whole to overspecialise along one particular dimension. However, pareto coevolution introduces some new problems: the pareto front is expensive to calculate; the size of the front may grow very large and be fairly unchanging; and it is not obvious how the final front should be combined as a solution set, i.e., which should be chosen when? In conjunction, the contrived nature of pareto coevolution makes it an unrealistic tool for biological modelling, thus restricting usage to purely problem-solving domains.

### 3.4.3 Cycling

#### In Nature

Ever since coevolution emerged as an independent discipline in the 1970s, cyclic coevolutionary trajectories have been anticipated (see Maynard Smith, 1982; Futuyma & Slatkin, 1983). Dawkins and Krebs (1979) suggest that (in asymmetric systems) cycling may be common both *between* and *within* species. They illustrate this by highlighting the genetic model of parent-offspring conflict postulated by Parker and Macnair (1979) and Parker (1979). A dominant “conflictor” gene causes offspring to demand more investment than the parental optimum. This is countered by the spread of a “suppressor” gene in parents

<sup>10</sup>Perhaps more accurately termed pareto hill climbing due to the elitist nature of the pareto front.

which causes them to invest equally in offspring, irrespective of demand. Assuming that the conflictor behaviour has a cost, then the direction of selection on children is reversed once suppressor genes are frequent amongst parents: non-conflictor genes spread and the cycle starts again.

Dawkins and Krebs (1979) also present an example in the context of predator-prey coevolution. Each individual pays a cost of size: predators pay the cost of growing big enough to swallow prey and prey pay the cost of growing big enough to prevent being swallowed. Hence, there ensues a coevolutionary race as each species' body size grows until their upper limit is reached; the size at which cost outweighs reward. Given the "life-dinner principle",<sup>11</sup> prey are likely to invest more in getting bigger (and thus have a higher limit) than predators, which are equally likely to divert scarce resources towards other adaptations (sexual attractiveness, for example). The side that can afford to pay the highest cost will do so; at which point selection will favour a rapid reduction in the cost paid by the other side. Any return to the start state may result (after some possibly unpredictable period of time) in a repeated bout of escalation. Maynard Smith (1996) describes such cycles as "sawtooth" oscillations and suggests that they can occur in general when a variable, such as size, can vary continuously with no ESS: hence, size may increase gradually until a threshold is reached, when the population can be invaded by much smaller individuals.

Sawtooth oscillations have been observed in *Anolis* lizard populations of the Caribbean islands (Roughgarden, 1983a). The insects on which *Anolis* lizards feed have food value proportionate to their size (the bigger the better, as long as they can be swallowed), and abundance inversely proportional to size (smaller insects are more common). A relationship between lizard size and insect size exists such that larger lizards on average take larger prey. On all islands inhabited by only one species, lizards have an equilibrium "solitary size" of approximately 50mm length in females and 60mm length in males. Where two species exist, however, one population is generally much larger than the solitary size whilst the other is *always* smaller. Where two species compete, it is assumed that the larger species exerts a stronger pressure on the smaller species than *vice versa*: large lizards take more food away; disputes over territory favour large lizards. To fit this data, Roughgarden (1983a) proposed the "coevolution-invasion turnover hypothesis", suggesting that through a series of invasions and extinctions, the lizard populations cycle in body length (see figure 3.3, adapted from Roughgarden, 1983a). On islands containing only one species, lizard lengths approximate the optimum solitary size; because of the asym-

---

<sup>11</sup>The life-dinner principle states that coevolving species may be subject to asymmetric evolutionary pressures. For instance, consider foxes and rabbits: "The rabbit runs faster than the fox, because the rabbit is running for his life while the fox is only running for his dinner" (Dawkins & Krebs, 1979, after Aesop).




Figure 3.3: Invasion-extinction cycles in *Anolis* lizards of the Caribbean islands. (A) Islands containing one resident population evolve to the equilibrium solitary size. (B) Lizard populations with larger body size are able to invade. Both populations evolve towards smaller body size. (C) The resident population becomes extinct, leaving the invaders to evolve to the solitary size. The cycle is complete (adapted from Roughgarden, 1983a).

metry in competition, only a larger species can invade (figure 3.3.A). After invasion, the species coevolve as competitors. The resident's body size reduces to avoid competition from the invader; an example of character displacement. The invader's body size also reduces to take advantage of the greater resources left by the retreating residents (figure 3.3.B). Finally, the resident species is driven to extinction by competitive exclusion. The invaders approach the optimum solitary size and the system has completed one full cycle (figure 3.3.C).

Since the example of the *Anolis* lizards requires species replacement (an intraspecific cycle of the same nature would require invasion by an implausibly large mutant), Maynard Smith (1996) suggests that the first example of a population cycle arising from intraspecific interactions is that reported relatively recently in side-blotched lizards (Sinervo & Lively, 1996). Side-blotched lizards exhibit three alternative male mating strategies, each associated with a distinctive phenotypic trait: blue-throated males mate-guard females and are territorial; yellow-throated “sneaker” males are non-territorial and roam about freely, looking to copulate with the females of others; and aggressive orange-throated males are polygynous and maintain large territories. Whilst blue-throated lizards avoid cuckoldry by yellow-throated sneakers, they are easily overpowered by orange-throated males, which cosire offspring with their females. Yellow-throated males are able to sire offspring via secretive copulations with the females of orange-throated males and often share paternity of offspring within a female's clutch (Zamudio & Sinervo, 2000). Sinervo and Lively (1996) showed that the frequencies of the three male morphs were found to oscillate over a six year period.

The fitnesses of each morph relative to other morphs were non-transitive in that each morph could invade another morph when rare, but was itself in-

vable by another morph when common. Concordance between frequency-dependent selection and the among-year changes in morph fitnesses suggest that male interactions drive a dynamic ‘rock-paper-scissors’ game (Sinervo & Lively, 1996).

Using quantitative measures of the reproductive success of males adopting each strategy, Zamudio and Sinervo (2000) confirmed that the morphs were indeed playing RPS. The relative fitness of each strategy during dyadic interactions confirmed this.

Intransitive RPS dynamics have also been observed in populations of the bacteria *Escherichia coli* (Kerr, Riley, Feldman, & Bohannan, 2002). Colicinogenic bacteria (C) possess a ‘col’ plasmid, a toxin that kills colicin-sensitive (S) bacteria. A third strain (R) is resistant to the colicinogenic bacteria. In some cases, the growth rate of R cells will exceed that of C cells, but be less than the growth rate of S cells.

In such cases, S can displace R (because S has a growth-rate advantage), R can displace C (because R has a growth-rate advantage) and C can displace S (because C kills S). That is, the C-S-R community satisfies a rock-paper-scissors relationship (Kerr et al., 2002).

Confirming the predictions of their simulation model, empirical observations of *E. coli* (constrained to local interactions in a petri dish) demonstrated the cyclical coexistence of all three strains, with R outperforming C, C outperforming S and (a suggestion of) S chasing R across the plate (refer to chapter 7 for further discussion of *E. coli*).

Using the public goods game (applicable to theoretical biology), cycling has also been demonstrated in experimental economics; both in theory (Hauert, De Monte, Hofbauer, & Sigmund, 2002) and in practice (Semmann, Krambeck, & Millinski, 2003).

Semmann et al. (2003) demonstrate that voluntary participation in the game can lead to cooperation amongst sizable groups; despite anonymity, random assortment and non-repetition of interactions. Three strategies exist within a population: defectors (D) and cooperators (C), both willing to engage in the public goods game and speculate on the success of a joint enterprise; and low-risk loners (L) who choose to reject participation and settle for a small, but guaranteed, return. From time to time, groups of individuals are offered the choice to compete in a public goods game; loners will always refuse.

In every group, defectors outperform cooperators, but if all cooperate, they are better off than if all defect. Whilst it is better to be a loner than in a group of defectors, it is better still to be in a group of cooperators. Hence, in a well-mixed population, strategies display a RPS cycle. If most play C, then it is better to play D, but if most play D, then it is better to play L. However, if most play L, then small groups can form, increasing the chance of

mutual cooperation. Thus, C dominates if group size is small, D dominates if group size is large, and the option to be a loner preserves a balance between the two options: the system cycles (Hauert et al., 2002). An empirical study involving 280 students playing a 57-round strategy game confirmed these results (Semmann et al., 2003).

### **In Artificial Systems**

Having been long predicted and observed in nature, it is unsurprising that cycling has received much attention in the artificial coevolutionary literature (e.g., Cliff & Miller, 1995; Rosin, 1997; Shapiro, 1998; Juillé & Pollack, 1998a; Ficici & Pollack, 1998b; Bucci & Pollack, 2002). Cycling occurs as a direct consequence of relative fitness: as individuals are only rewarded for out-performing *contemporary* opponents, it is possible for earlier adaptations to be lost through time. The resulting lack of evolutionary “momentum” may lead to cycling. Other effects, such as intransitive superiority relationships (i.e., rock-paper-scissors) and repeated arms races and crashes (like that described by Dawkins & Krebs, 1979), can also trap a system in a repetitive cycle. As the system stops progressing in any objective sense, the result can be disastrous for problem-solving.

Forming a direct lineage back to the evolutionary selection method of elitism (originally De Jong, 1975, see section 2.2.4), many techniques have been developed to combat coevolutionary cycling; each an attempt to install evolutionary memory and momentum. Elitism preserves the best individuals within a population by copying them *unchanged* into the next generation. Whilst not specifically developed for cycling (or for coevolution) elitism has since been adapted as a coevolutionary technique. The first to utilise elitism for relative fitness assessment was Sims (1994a). When coevolving 3-dimensional morphologies and behaviours, Sims chose to assess each individual against the same opponent, thus reducing sampling error (a potential hazard of random opponents) whilst still reducing computational cost (one assessment is cheaper than many). The fittest individual from the preceding generation was chosen for assessment each time. This installed some extra coevolutionary “direction”, with individuals forced to perform well against the same opponents that outperformed many of their parents.

Sims’ approach was repeated successfully by Reynolds (1994) and again by Cliff and Miller (1995), who referred to it as LEO (Last Elite Opponent) evaluation: using a two-population system, Cliff and Miller (1995) evaluated each individual against the previous generation elite from the *opponent* population. Floreano and Nolfi (1997b) increased the memory and momentum (and hence stability) of LEO by testing each individual against the best competitors of the previous 10 generations.

Further extending LEO contests to their ultimate extreme, Rosin (1997) produced the *Hall of Fame* (HoF). Whereas Cliff and Miller (1995) had collected information on rel-

ative success across generations and presented it to the observer (in the form of CIAO plots), Rosin's (1997) coevolutionary algorithm used the same information (in the form of a Hall of Fame) to guide its assessment of individuals. Every generation the best individual from each population is retained in the HoF. Each successive generation, individuals compete not only against a sample of opponents from the coevolving population, but also against a sample from the HoF. In this way individuals are selected to not only out-compete current opponents, but also to out-compete ancestral opponents. The idea is to provide a coevolutionary *memory* in order force progress in a particular direction and thus stop cycling. However, unlike the LEO predecessors, HoF is computationally expensive.

Rosin's HoF was extended by Olsson (2001) with the Asymmetry Handling Host-Parasite Algorithm (AHHPA). Olsson's extension aimed to improve computational efficiency. Rather than add the best individual from each generation to the HoF, a host, for example, is only added to the host HoF if it can beat all the current individuals in the parasite population *and* all the individuals in the parasite HoF. Hosts evolve against a static parasite population until such a successful host is found. This host is added to the host HoF and the host population is frozen whilst the parasites evolve until one is found that can beat every host in both the current and HoF populations. Thus, the HoF only grows each time an *unbeatable* individual is found. As such, it is much smaller (and less computationally expensive) than Rosin's Hall of Fame. In conjunction with computational efficiency, the AHHPA is intended to reduce redundant assessment due to asymmetry. Suppose that a coevolutionary system is much easier for parasites than hosts. For instance, this could be because parasites have much smaller genomes, thus making genotype space smaller and solutions easier to find. With Rosin's Hall of Fame method hosts and parasites evolve at the same rate. However, if hosts are finding it harder than parasites, they may benefit from more generations of evolution than their parasitic competitors. AHHPA reduces the number of *wasted* generations by allowing the host (or parasite) population to have more evolutionary generations than parasites (hosts). Thus, in the face of asymmetry, AHHPA allows coevolution to use computational time more effectively.

The Hall of Fame has been demonstrated to improve optimisation in several domains (Rosin, 1997; Rosin & Belew, 1997), however, a serious problem exists. Over time, the HoF grows with respect to the contemporary population of evolving individuals. As a result, fitness assessment becomes relatively more static through time: as a proportion of the HoF, the number of reciprocally evolving opponents tends to 0. Thus, as a run progresses, coevolution with the HoF tends towards standard evolution (assessment against a non-evolving function).

... as the process goes on, there is less and less pressure to discover strategies that are effective against the opponent of the current generation and greater and greater pressure to develop solutions capable of improving performance against opponents of previous generations (Nolfi & Floreano, 1998).

The memory HoF adds to coevolution becomes so great that not only are cycles prevented, but relative fitness ceases altogether. At such times, the coevolutionary system is better represented by evolution; with the problem of local optima re-emerging (for a similar argument see Luke & Wiegand, 2002).

### 3.4.4 Mediocre Stability

During a coevolutionary arms race, it is possible for the system to stabilise at a sub-optimal equilibrium: a problem analogous to convergence at a local optimum in evolution (see Angeline, 1994; Pollack et al., 1996; Juillé & Pollack, 1998a; Ficici & Pollack, 1998a; Pollack & Blair, 1998; Bucci & Pollack, 2002). Interspecific *collusion* is often cited as the cause of mediocre stability. This can be highlighted with a simple example. Imagine two chess-playing populations are coevolving. Individual chess strategies are assessed through a series of matches, with each taking alternate turns to play white. In this scenario, individuals can improve their chances of reproduction in two ways. Firstly, they can play as well as possible in order to win and receive high fitness. Secondly, they can throw every game they play black: if an opponent does the same, collusion guarantees both 50% of victories. Whilst the first scenario leads to progress, the second results in mediocre stability that can be difficult to escape.

Intraspecific collusion is also possible. Imagine a population of coin toss predictors is coevolving against a population of coin toss generators. The required optimum is a random set of generators (50% heads, 50% tails) and a population of predictors guessing each outcome with equal probability. However, a mediocre stable state also exists. At the global optimum, each individual is, on average, winning 50% of competitions. But, this state can also be reached if the predictors diversify into two separate niches, one always guessing the outcome heads and the other always guessing outcome tails. At this mediocre stable state, the generators and predictors still score 50% each, but the outcome is undesirable as the coin toss predictors are each highly brittle.<sup>12</sup> Unfortunately, escaping

---

<sup>12</sup>Whilst individuals within each niche are themselves brittle, the population as a whole *does* contain information necessary to produce generalist strategies. As Yao, Liu, and Darwen (1996) discuss, when using (co)evolution for learning (rather than optimisation), information from the whole population should be retained. Integrating these strategies, however, is far from trivial.

from this sub-optimal equilibrium may be difficult (for further discussion of these issues see section 6.2).

A famous example of collusion occurred in the trenches of the Western Front during World War I. Whilst locked in stalemate, Allied and German troops both refrained from killing enemy soldiers unnecessarily. Rituals were adopted by both sides: artillery was fired at the same time each day and mealtimes were particularly avoided. Collusion was only ended when officers demanded a series of changes: troops were regularly rotated, and small-party raids into enemy trenches were ordered (Axelrod, 1984).

Generally, mediocre-stability is difficult to classify in biology. As previously discussed (refer to 3.2.4), one cannot determine objectively whether a biological system is at a sub-optimal equilibrium or not.

As far as the author is aware, no technique has been specifically developed to combat mediocre stability. Amongst the problems of coevolution (discussed here), this makes mediocre stability an exception. Under some circumstances, however, other techniques may combat mediocre stability as a beneficial side-effect.

### 3.4.5 Disengagement

Coevolving populations can be considered as a coupled dynamical system with each population evolving over a dynamic fitness landscape that continually fluctuates in response to perturbations from the other (see discussion of *NKC* landscapes in Kauffman & Johnsen, 1991). It is this *interactive* dynamic that drives the selection pressure between populations, continually eroding the adaptive advantage of each population (the Red Queen Effect, van Valen, 1973) and potentially resulting in an evolutionary arms race. Disengagement occurs when a coevolutionary system *decouples*—each population no longer perturbs the other, thus not only eliminating feedback between the coevolving populations, but eradicating any means of relative fitness assessment—resulting in evolutionary drift (for a more extensive definition of disengagement, refer to 3.4.5.1).

Disengagement occurs when one population *outcompetes* the other to the extent that different individuals are not discriminated from their contemporaries in fitness terms, i.e., floor or ceiling effects (e.g., Watson & Pollack, 2001). Often, coevolutionary systems are asymmetric—hosts and parasites may differ genetically (in terms of encoding) or behaviourally (in terms of goal strategy). Such asymmetry may result in an inherent advantage for one population. When coevolving pursuers and evaders, for example, it is often much easier, at least initially, to be a successful evader (Cliff & Miller, 1995). Given that disengagement results from one population outcompeting the other, it is intuitive



that an inherent asymmetrical advantage favouring a particular population will encourage coevolutionary disengagement.

Consider an asymmetrical host-parasite system in which parasites enjoy an inherent advantage. Let us assume that the system is nearing disengagement, with the majority of parasites scoring maximally in the majority of competitions against hosts. The few parasites that some hosts are able to beat—those that discriminate hosts—will receive relatively low fitness and as such will have few progeny. In contrast, the parasites able to beat all current opponents—and thus unable to discriminate among hosts—receive high fitness, thus leaving many offspring. In such a situation, it is likely that subsequent parasite generations will tend to comprise increasing numbers of individuals capable of beating all current opponents—i.e., there will be less discrimination among hosts despite any genetic and phenotypic diversity. Eventually the populations will disengage, with every host achieving the same poor score, and every parasite achieving the same high score. At this point, both populations will drift.

Although disengagement has entered the terminology of coevolutionary computing relatively recently (Bucci & Pollack, 2002; Cartlidge & Bullock, 2002), the phenomenon *has* previously been recognised. Coevolutionary coupling, or engagement has been described as maintaining a gradient for selection with which to discriminate individuals (Watson & Pollack, 2001), coevolving an ideal training set with which to supply feedback (Juillé & Pollack, 1998a, 1998b), maintaining learnability (Ficici & Pollack, 1998a, 1998b) or providing pedagogical stepping stones (Rosin, 1997). The decoupling of coevolutionary populations is problematic and as such there have been several attempts to produce a technique to circumvent it. Methods proposed to counter-act disengagement are discussed below.

The *phantom parasite* (Rosin, 1997) inhibits the reproductive influence of unbeatable individuals (figure 3.4, left). Used in conjunction with competitive fitness sharing (Rosin & Belew, 1995), this *anti-elitist* innovation has no effect on individuals who achieve a less than perfect score. However, any individual that scores perfectly automatically has its score reduced slightly. This punishment is glossed in terms of interactions with an “ideal” phantom parasite: ‘Hosts that lose to some current parasite defeat the phantom parasite. Hosts that defeat all current parasites lose to the phantom parasite’ (Rosin, 1997). Effectively, the phantom parasite transforms the fitness function of parasites as shown in figure 3.4, left. Individuals winning  $N - 1$  contests receive an equivalent score to those winning  $N$  contests. The aim is to discourage parasites from performing *too well*, however, it remains in the *best interests* of an individual parasite to win as many competitions as possible.

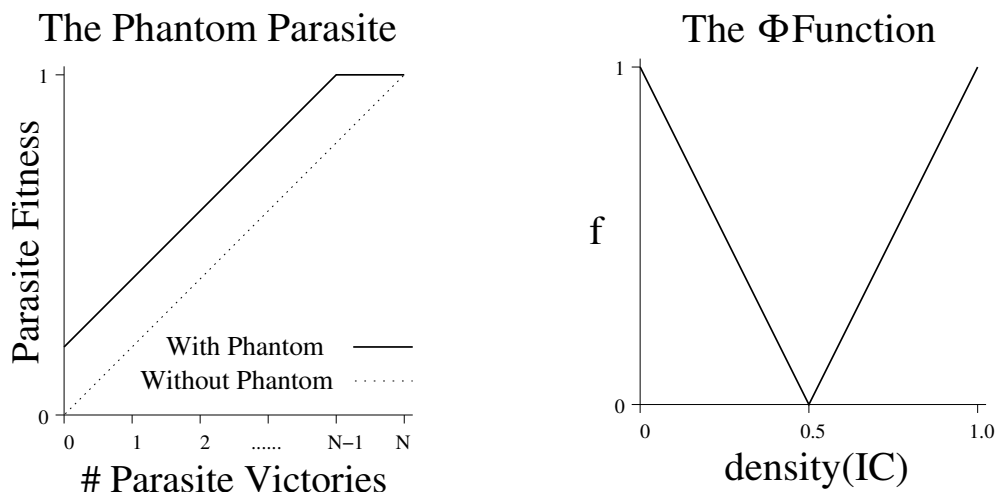


Figure 3.4: The phantom parasite function (left). Individuals that beat all *real* opponents lose to the phantom parasite. Individuals that lose to at least one real opponent beat the phantom parasite. The  $\Phi$  function (right)—fitness of an IC is 0 if classified correctly,  $f$  otherwise—produces a *stabilising* selection pressure favouring ICs that are neither too easy, nor too hard, to classify (see equation 3.1).

The *friendly competitor* (Ficici & Pollack, 1998b, 1998a) attempts to coerce a coevolutionary arms race through utilising a three population system. A parasite population coevolves with two host populations, one *friendly* and one *hostile*. Parasites are rewarded for being an easy challenge for friendly hosts whilst simultaneously being a difficult challenge for hostile hosts. When coevolving sorting networks, for example, a parasite list would be rewarded for simultaneously being unsorted by hostile host networks and sorted by friendly host networks. In this way, parasites are punished for becoming too difficult for all hosts. The friendly competitor is designed to maintain a selection gradient for hosts, however, it is a complicated technique that requires the introduction of a third population and assumes that it is possible for parasites to simultaneously be both competitive and cooperative. To enable versatility, a simpler approach is required, such as that of the entropy measure, discussed below.

The *entropy measure* (Juillé & Pollack, 1998a, 1998b), later adapted as Pagie’s  $\Phi$  function (Pagie & Hogeweg, 2000; Pagie & Mitchell, 2002), utilises domain-specific knowledge for the density classification task for cellular automata (CAs) to reduce disengagement in that specific domain (figure 3.4, right). The aim of the density classification task is to coevolve CA rules for classifying the density of an initial condition (IC)—a binary string—as greater (or less) than 0.5 depending upon whether the IC contains more (or less) than 50% 1s (see Mitchell, Crutchfield, & Hraber, 1994). Having an inherent advantage, parasite initial conditions are known to become increasingly difficult to classify

as their *density* approaches 0.5. Thus, in order to encourage ICs to be challenging, without becoming too difficult, engagement is maintained by utilising the  $\Phi$  function, where  $\Phi(IC)$  gives the fitness of an initial condition (Pagie & Mitchell, 2002).

$$\Phi(IC) = \begin{cases} 0 & \text{if classified correctly} \\ |density(IC) - \frac{1}{2}| & \text{otherwise} \end{cases} \quad (3.1)$$

As parasite ICs are encouraged to be unclassifiable, it is likely that IC density will approach 50%—the most difficult to classify. However, the  $\Phi$  function counteracts this by simultaneously rewarding ICs for deviating from 50% density. Thus, optimal ICs are as *easy as possible* to classify whilst still being unclassifiable. The  $\Phi$  function performs well, but is domain-specific. Although Juillé and Pollack (1998a) state that they would like to produce heuristics to make this technique domain-general, both the entropy measure and the  $\Phi$  function rely heavily upon domain-specific knowledge.

In summary, although disengagement is a recognised hindrance to coevolution, few techniques have been proposed to counteract it. Those that have suffer from problems such as domain-specificity or lack of versatility. The only real domain-general solution that has been proposed is the Phantom Parasite, however, this has rarely been used (also, see chapter 6).

### 3.4.5.1 Classifying Disengagement

To the author’s knowledge, while there exists a general conceptual understanding of disengagement, no rigorous classification exists. To rectify this, the following classifications—to be used throughout this thesis—are proposed (also Cartlidge & Bullock, 2004a).

*Full disengagement* occurs when selection cannot distinguish between individuals irrespective of which sample of the current opponent population each individual plays. When assessment is noisy, although individuals in disengaged populations will achieve different scores, selection will not be able to distinguish between them on any systematic (i.e., non-random) grounds. In this situation, a coevolutionary algorithm has no basis upon which to preferentially select certain genomes for reproduction. Coevolutionary drift ensues, allowing deleterious mutations to accumulate in each population. Typically, in the course of running a coevolutionary algorithm, episodes of full disengagement will be difficult to distinguish from a much weaker kind of “contingent” disengagement. In this case, although each member of a population achieves identical, or near-identical fitness scores, they *could* have achieved different scores had they played alternative members of the current opponent population. We will use the term disengagement to refer to the complete lack of selective discrimination that results. The *degree* of disengagement, mea-

sured by how far (in terms of novel mutations) the populations are away from re-coupling, allows us to determine the likelihood of a re-engagement event.

*Asymmetrical disengagement* occurs when one population,  $P_1$ , reaches the global optimum whilst the other,  $P_2$ , drifts through sub-optimal space, resulting in asymmetric selection pressure: whilst  $P_2$  drifts, sub-optimal mutants from  $P_1$  are selectively punished. An example of asymmetrical disengagement can be observed in the coevolutionary optimisation of sorting networks and test-lists. Once the sorting networks reach optimality, the selection pressure upon test-lists disappears—as all lists are sortable, each is equivalent.<sup>13</sup> However, the sorting networks are under pressure to remain at optimality as sub-optimal mutants may be punished by test-lists<sup>14</sup>. Alternatively, asymmetrical disengagement may be observed in the coevolution of poker players. Players giving away information for free—via either verbal or physical signals—are disadvantaged. As opponents gain from correctly interpreting the signals that they receive, the optimal strategy for a player is to not signal at all. However, once a player stops signalling, all receiver strategies become equally redundant—there is no benefit in attempting to read any cues. Asymmetrical disengagement occurs when players reach the non-signalling optimum. Whilst receivers drift, any signaller deviation—i.e., starting to signal—may be punished.

*Collusive disengagement*—similar to mediocre stability (section 3.4.4)—occurs when intra-population collusion produces inter-population disengagement. For example, once again consider coevolving populations of coin toss generators and coin toss predictors. Collusion between predictors is possible: by dividing into two niches, one always predicting “heads” and the other always predicting “tails” (assuming generators are unbiased) each predictor will be correct approximately half the time. As a result, predictors and generators become indistinguishable through fitness alone—there is disengagement. However, unlike full or asymmetrical disengagement, neither population is free to drift; since mutations away from each niche are punishable the *status quo* is likely to be preserved (refer to section 6.2).

---

<sup>13</sup>Whilst asymmetrical disengagement—the scenario of “perfect sorters”—is desirable here, it will not always be so.

<sup>14</sup>Asymmetrical disengagement may rapidly become full disengagement. Considering sorting network coevolution, the drifting list population is likely to become *easier* to sort in some absolute sense, thus reducing the selective pressure upon networks and enabling the accumulation of deleterious mutations.

In biology, when disengagement occurs in pairwise coevolutionary systems (e.g., symbiosis), at least one population is likely to go extinct. If there is a strong mutual dependence, both species may be driven extinct. Consider the interaction between myxoma and rabbits. If the myxoma virus ever becomes so deadly that rabbits are driven extinct, having no alternative host, the myxoma population is destined to follow.

However, disengagement does not necessarily result in extinction. Using the geographic mosaic theory (Thompson, 1994, also refer to section 3.2.1), it may be possible for populations to disengage locally but remain engaged (perhaps with other species) across their geographic range. Evidence for disengagement in biology can be observed in the invasion-extinction cycles of *Anolis* lizards (section 3.4.3; figure 3.3). Disengagement occurs each time a resident population goes extinct.

The fossil records of large cursorial predators (those that pursue prey for long distances over flat terrain) and their cursorial prey (usually ungulates, which respond to attack by fleeing at high speed) have been analysed (Bakker, 1983). Such analysis is possible because the functional anatomy of running on flat surfaces is well understood. Since most ungulate species rely entirely on flight for defence, the expected coevolutionary response from prey (with deer-like body forms) should be toward the longest, swiftest limb form possible. Likewise, predators should follow a parallel trend for increasing speed. However, over much of the last 60 million years, predators have lagged behind in evolutionary terms. Predators and prey have *not* evolved at the same pace—there is an “evolutionary gap”. This suggests some degree of disengagement. In the multi-lineage coevolution between cursorial predators and ungulate prey, prey are generally winners.

Rather than corroborating the notion of quick optimisation and congruent coevolution, the fossil record of distance-pursuit predators shows that predators fell behind their prey in the arms race of locomotor advancement (Bakker, 1983).

### 3.4.5.2 Coevolutionary Neutrality

Search-space neutrality occurs when many genotypes share the same fitness, perhaps as a result of redundancy in the genetic encoding. A neutral set contains all the genotypes that achieve a particular fitness score, while a neutral network comprises evolutionarily adjacent members of a neutral set. It has been argued that the presence of neutral networks may have profound consequences for the dynamics of evolutionary search. For instance, the neutrality exhibited by natural RNA search spaces has been demonstrated to be of

a potentially useful kind, allowing more efficient search (Huynen, Stadler, & Fontana, 1996). More generally, neutrality of the right kind is thought to reduce the chance of premature convergence (Harvey & Thompson, 1996; Barnett, 1998).

In a coevolutionary system, since the fitness scores of individuals are calculated relative to reciprocally evolving individuals, the topology of neutrality is continually shifting: genotypes may be neutral in one environment, but not another. During full disengagement, since individuals cannot be discriminated, each population forms a neutral sub-set of equivalent fitness. The system drifts until re-engagement occurs. However, unlike “useful” neutral-drift, where the “quality” of genotypes is preserved (i.e., genotypes are equivalent in some meaningful sense), drift during disengagement is effectively random and thus potentially degrading. Often, neutrality resulting from disengagement is *not* “useful”.

For an in-depth treatment of disengagement and the related phenomenon of coevolutionary neutrality, refer to chapter 8.

### 3.5 Classroom Analogies

To exemplify the problems discussed in the previous section, coevolutionary theory is applied to trends in exam results in England and Wales over the last three decades.

**Visualization issues.** A-level pass-rates have risen for 21 consecutive years (Harrison, 2003). Depending upon political bias, this can *either* suggest that students are getting progressively better (a victory for the education system) or that exams are getting progressively easier (a dumbing down of the A-level “gold-standard”). In truth, the relative nature of assessment—exams adapt in response to changes in the teaching syllabus—unsubstantiates both claims on the basis of this data alone. For the same reasons that simple coevolutionary “fitness over time” graphs are moot, comparing A-level results over time is uninformative. To investigate whether students have improved in real terms, it is necessary to use other methods of comparison. Perhaps a fixed (and secret) “benchmark” exam could be taken each year by a small subset of A-level students. More objective comparison between these results and the A-levels students subsequently obtain would then be possible. For example, if benchmark results were shown to remain fairly constant whilst A-level results continue to rise, then a strong case for “dumbing down” could be presented.

**Overspecialisation.** The education department is under continuous political pressure to improve the average exam performance of students. Since the *raison d’être* of ex-

ams is to indicate a student's all-round ability in a given subject, a general improvement in exam grades is often assumed a reflection of improving education standards. However, it is possible for students to improve exam results by over-specialising upon the specific exams they will face. In such circumstances, exam results lose accuracy as an indicator of general ability: students are good at passing exams, but not much else. It has recently been suggested that "teaching to the test" may be exacerbating overspecialisation, with students trading general ability for specific results. Specialization has grown as a result of national targets: 'If it appeared in the tests, then it was important. If it didn't, it was a luxury' (Baker, 2002a).

**Cycling.** Directed by government policy, education strategy (and related testing) continually adapts to fulfil the political requirements of the nation. Imagine there is a dearth in European language skills, for example. As a result, education funds are likely to be channelled into modern languages, perhaps to the detriment of other subjects, such as the sciences, say. Before long, the adapting student population will possess more language skills, but fewer scientists. The change in political demand reverts education strategy, returning to a focus on science. Language skills decline and the system cycles.

Such cycling can be observed in England and Wales over the past two decades. Having been biased towards the three Rs (reading, writing and arithmetic) for many years, education policy was radically re-worked in 1989 with the introduction of the national curriculum. Demanding a broad range of ten discrete subjects, the national curriculum was designed to give creative, practical and academic skills to students. However, employers' protests over falling standards in the "key-skills" has led to a more recent drive for literacy and numeracy. Mike Tomlinson, the chief inspector for Ofsted (Office for Standards in Education) suggests that in half of all schools, English has now commandeered between a quarter and a third of all lesson time (Baker, 2002a).

**Mediocre stability.** In 1995, the introduction of SATs (standard assessment tests) heralded the much anticipated arrival of school league tables. Forcing schools into direct competition, it was hoped that the publication of SAT results would rejuvenate the education system and allow poor performers to be targeted. However, it seems much more likely that a state of collusion has been reached. Unlike GCSEs and A-levels, SATs are much more of a performance indicator for teachers, rather than students—if SAT results are poor, a school's league table position may directly

affect the job security of its teaching staff. Thus, to secure their job, a teacher has two choices: work harder for no extra pay and no guarantee of reward; or cheat. Given that other teachers are in the same dilemma, game theory (multi-player prisoners dilemma) suggests that defection (cheating) will be widespread. Teachers are likely to collude with students, influencing the results of their SATs through the use of visual cues, coaching, or direct alteration. In 2001, 11 primary schools had their results annulled after investigations revealed that they could not be sure the pupils' work was their own. One head teacher admitted altering answers (Baker, 2002b).

**Disengagement.** Universities discriminate between students on the basis of their A-level results: if entry grades are satisfied, a student is usually offered a place. For this process to succeed, A-level grades must be informative—there must be a spread of results within the student population. However, in recent years, A-levels exams and students have begun to disengage. With the continual rise in grades, many students are now outperforming the exam. In August 2003, 38.9% of mathematics students received an A grade at A-level (Harrison, 2003). As a result, it has become increasingly difficult for universities to differentiate between the best students. Some institutions are considering the revival of entrance exams.

### 3.6 Summary

A system is classified “coevolutionary” if and only if each evolving individual is assessed against a reciprocally evolving function. The dynamic fitness landscape that results can circumvent some of the traditional problems of non-coevolutionary (or standard evolutionary) algorithms, such as premature convergence at a local-optimum. As a result, coevolution has been successfully applied in many problem domains. However, the relative fitness “seascape” of coevolution also introduces several new problems: visualization; overspecialisation; cycling; mediocre-stability; and disengagement. Each of these can be observed in the real-world coevolution between students and exams.

Throughout this thesis the problems of coevolution are teased apart and considered in isolation, thus allowing experimental factors to be more easily isolated whilst making empirical analysis more tractable. However, it may be objected that, if the problems of coevolution are highly inter-related, then the process of understanding and tackling them in isolation is potentially futile. Yet, because of the immaturity of research in this area—this thesis is perhaps the first extended attempt to solve one of coevolution's problems—it is too early to know whether or not this methodological gambit is a high risk strategy.



Hence, taking the only sensible approach available, it is presumed that, for the sake of argument, teasing apart the problems of coevolution will afford an eventual synthesis of knowledge.

# Chapter 4

## Research Statement

---

### 4.1 Objective

To improve the application of coevolutionary algorithms to complex problem solving, this thesis pursues two simultaneous goals: (1) to further theoretical understanding of coevolutionary algorithms; (2) to propose novel techniques for improving the performance of coevolutionary algorithms. More specifically, the following sub-goals are pursued:

#### Definitions

- To rigorously classify coevolution in both biological and artificial systems, allowing theory crossover between evolutionary biology and evolutionary computation (already developed; chapter 3).
- To classify coevolutionary disengagement (already developed; chapter 3).

#### Disengagement

- Investigate disengagement in a variety of coevolutionary systems. Further understanding of the causes and ultimate effects of disengagement.
- Develop a novel technique capable of combating coevolutionary disengagement.
- Investigate the effects of this novel technique. Compare performance with current state-of-the-art approaches.

## Visualization

- Investigate the suitability of current coevolutionary visualization techniques.
- Propose alternative/novel visualizations capable of furthering the theoretical understanding of coevolutionary dynamics.

## Neutrality

- Highlight the role of neutrality in coevolutionary systems, emphasising the relationship between neutrality and disengagement.
- Outline a novel framework for investigating coevolutionary neutrality.

## 4.2 Justification

In the late 1970s it became clear that, without a universally accepted definition, the term “coevolution” was in danger of losing practical significance in evolutionary biology (refer to section 3.2). Today, a similar situation faces evolutionary computing. Whilst “coevolution” is often discussed, no rigorous classification exists. The development of a new classification scheme capable of encompassing both biological and artificial systems would not only offer a strict theoretical framework within which to study coevolution, but would also offer potential crossover between evolutionary biology and evolutionary computing (underpinning this thesis, a novel classification of coevolution has already been introduced; section 3.2). Similarly, coevolutionary disengagement is a phenomenon that is often hinted at, but that has no formal definition. As before, a strict definition is required before disengagement can be adequately studied (again, previously introduced; section 3.4.5.1).

Although several authors have observed the detrimental effects of disengagement, the phenomenon has never been studied in detail (section 3.4.5). Though techniques have been proposed to combat disengagement, they are often domain dependent and poorly tested. In general, work in this area is lacking. Based on a firm theoretical understanding of the causes and effects of disengagement, a novel technique capable of counteracting disengagement is needed.

To gain insight into the dynamics of coevolutionary systems, appropriate visualizations are crucial. However, since coevolution renders standard visualizations—such as simple graphs of fitness over time—redundant, novel techniques are required (refer to section 3.4.1). CIAO plots are perhaps the only visualization specifically developed for

coevolutionary systems and, as such, they are widely used. However, having received little critical investigation, the limitations of CIAO plots are not well understood: this situation must be rectified. Alternative visualization tools, offering greater insight into coevolutionary behaviour, are required.

In general, no distinction is made between coevolutionary neutrality and neutrality in non-coevolutionary systems (such as RNA folding landscapes; section 3.4.5.2). However, there is a fundamental difference: coevolutionary neutrality can vary through time and is intrinsically linked to disengagement. To fully understand disengagement—in particular, *why* it is that disengagement is detrimental—it is necessary to consider coevolutionary neutrality. As yet, there has been little work in this area; indeed, there is currently no available framework within which to study coevolutionary neutrality.

### 4.3 Prospects

In the natural world, since no species has evolved in complete isolation, coevolution is responsible for much of the complexity observed in plants and animals. To be able to harness a fraction of coevolution's power using coevolutionary algorithms is a tremendous prospect, with barely imaginable potential for the future of intelligent systems. Yet, despite a number of problems (section 3.4), coevolutionary algorithms have *already* achieved success in a number of problem domains (section 3.3). If these problems—visualization (3.4.1), overspecialisation (3.4.2), cycling (3.4.3), mediocre stability (3.4.4) and disengagement (3.4.5)—were better understood (and ultimately solved), the full potential of artificial coevolution may start to become a more realistic prospect; in conjunction, novel insights may also shed light on the behaviour of natural coevolutionary systems. Thus, fulfilling the objectives of this thesis (section 4.1) has the potential to further not only evolutionary computing, but evolutionary biology as well.

# Chapter 5

## Reduced Parasite Virulence

---

Designed to encourage arms race dynamics, conventional coevolutionary algorithms reward the “best performing” individuals of each generation with reproductive success. With each new innovation expected to outperform the last, it is generally accepted that a series of adaptations and counter-adaptations will ensue.

However, in this chapter, it is argued that the conventional approach to selection is slightly misguided: on occasion, the best way to encourage arms races may be to *punish* successful individuals; not reward them. Inspired by the virulence of natural parasites, “reduced parasite virulence” is developed; a novel technique designed to improve the performance of competitive coevolutionary algorithms by prolonging arms race engagement.

### 5.1 Rationale

#### 5.1.1 Here Today, Gone Tomorrow

Escalating arms races (Dawkins & Krebs, 1979) are the Holy Grail of artificial coevolution. With competing populations continually striving to gain an upper hand, arms race dynamics are expected to compel individuals towards ever more refined adaptations; resulting in accelerated evolution and, ultimately, a population of sophisticated solutions.

To induce an arms race, individuals best able to outperform opponents are selectively biased in favour of reproduction. The successful adaptations contained within these individuals are preserved in their progeny. Unsuccessful adaptations are discarded—“losing”

individuals are unlikely to reproduce. Hence, over evolutionary time, each population becomes better able to exploit the idiosyncratic weaknesses of the other. However, whilst this approach is designed to encourage arms races, the following problems threaten to destroy them.

**Thrashing.** Under the conventional coevolutionary approach, genes conferring the greatest advantage over contemporary opponents are those most likely to propagate. Conversely, the very opponents that successful individuals are able to beat are unlikely to be represented in proceeding generations. If individuals are highly specialised (perhaps *overspecialised*; section 3.4.2), successful adaptations in one generation may be effectively “worthless” in the next: particularly if a strategy is only successful against the competitors it has driven extinct. As a result, the system may “thrash about”; repeatedly shifting evolutionary direction. Saw-tooth oscillations may follow.

**Disengagement.** Exploiting the weaknesses of opponents may also lead to disengagement (section 3.4.5). When an individual in population  $A$  discovers an adaptation,  $\alpha$ , capable of exploiting individuals in population  $B$ , the successful adaptation quickly spreads. If every member of  $A$  inherits  $\alpha$  before a suitable counter-adaptation is discovered by population  $B$ , then disengagement will occur. All the members of population  $A$  (thanks to  $\alpha$ ) are able to beat all the members of population  $B$ . Individuals can no longer be discriminated, selection pressures disappear and the system drifts.

Since exploitable opponents in one generation may be absent from the next, biased selection in favour of the best individuals may *discourage* coevolutionary arms races. To *encourage* an arms race, both populations must contain adaptations capable of beating not only today’s opponents, but tomorrow’s opponents as well.

### 5.1.2 Don’t Bite the Hand that Feeds You

Coevolving populations are mutually inter-dependent. Since opponents are relied upon for success, even in competitive systems coevolution is not entirely antagonistic (sum-zero). The conventional selection methodology should reflect this dependence; however, it does not. To reduce the likelihood of thrashing and disengagement, selection should encourage the preservation of *some* exploitable strategies.

Here, a novel selection method (“Reduced parasite virulence”) is proposed: individuals are *only* rewarded for moderate victories—excessive victories are punished.<sup>1</sup> Thus, rather than succumb to the temptation of “biting the hand that feeds them”, individuals are encouraged to adopt a more Machiavellian approach to competition:

...in most cases, so long as you do not deprive them of either their property or their honour, the majority of men live happily; and you have only to deal with the ambition of a few (Machiavelli, 1998, p61).

In many political, economic and biological systems, an analogous trade-off exists—short term gains may be relinquished for long term advantage. The following section details some examples.

## 5.2 Analogies

### 5.2.1 Parasite Virulence

Following Janzen (1980; also refer to section 3.2.1), one can distinguish between *true* coevolution, consisting of two populations, and *diffuse* coevolution, consisting of multiple populations, each reciprocally evolving. Often pairwise, artificial coevolution is most accurately categorised as true coevolution. In particular, the mutual inter-dependence between artificially evolving populations strongly resembles that between symbiotic partners (especially hosts and parasites; given the antagonistic nature of competition).

[When considering symbiosis, because] the individuals of different species are physically associated with one another for a long time, the survival of at least one member of the association depends upon the survival of the other member. Hence, the fitness of an individual depends on the *fitness* of its associate. (Roughgarden, 1983b, original italics).

Thus, the best biological analogy of artificial (two-population competitive) coevolution is host-parasite coevolution—one population (the parasite) is typically considered to pose problems for the other (the host). However, in nature, direct evidence of host-parasite coevolution is exceptional. Rare examples include the tight step-for-step coevolution observed between the myxoma virus and rabbits (Fenner & Ratcliffe, 1965, also

---

<sup>1</sup>Upon first reflection, it may appear that the introduction of “virulence” in artificial systems is an unnatural addition to the coevolutionary algorithm. However, utilising an analogy of parasite virulence is no less realistic than Hillis’ (1990) original model. *Nevertheless*, virulence is not introduced to add realism, but to improve algorithm performance.

section 5.2.1.1); and the appearance of strains of the Hessian fly (*Mayetiola destructor*) that are able to attack a series of sequentially planted strains of wheat (Gallun, 1977).<sup>2</sup>

Coevolutionary algorithms typically differ from natural systems in the way that they deal with parasite virulence (here defined as ‘parasite-mediated morbidity and mortality in infected hosts’; Levin, 1996). In order to ensure survival long enough to reproduce, it is not always in the best interests of a biological parasite to be as virulent as possible (Futuyma & Slatkin, 1983; Maynard Smith, 1989; Combes, 1991; Maley, 1994; Hood, 1997). As a result, virulence varies dramatically between systems (compare, for instance, cholera and the common cold), and over time within a particular system (e.g., the history of the myxoma virus in Australian rabbit populations, Fenner & Ratcliffe, 1965, also section 5.2.1.1). However, when parasites are used in artificial coevolution, they are generally encoded to be maximally virulent—their fitness varies inversely with the success of the hosts that they compete against.

Whilst a parasite gains from using a host’s resources, incapacitating a host can greatly reduce a parasite’s chances of reproducing. For this reason, a parasite’s virulence and dispersal distance are strongly linked (Maley, 1994). A parasite that greatly relies on the mobility of a particular host is likely to have relatively mild virulence. In contrast, parasites that are easily transmitted from one host to another, irrespective of host mobility, are likely to remain highly virulent. Using mosquitoes as a transmission vector, malaria can afford to fatally damage hosts without hindering propagation. Indeed, a host rendered unconscious is more susceptible to mosquito bites, ultimately aiding transmission. In stark contrast, as the common cold spreads via host-host contact, it benefits from keeping hosts active and sociable.

The virulence of natural parasites strongly affects the coupling between host and parasite populations. Extremely virulent parasites may ultimately push their hosts—and consequently themselves—to extinction, resulting in a decoupling of the system. As persistent high virulence can result in the disengagement of natural host-parasite systems, we enquire as to whether reduced parasite virulence could reduce the effects of disengagement in artificial coevolutionary systems—might coevolutionary algorithms benefit from treating parasites more naturally?

In the following section, myxomatosis-rabbit coevolution is detailed to illustrate the dynamics of parasite virulence in biological systems. It is demonstrated that the “optimum” (evolutionarily stable strategy, or ESS) value of virulence depends upon the intricate nature of the system. Unless otherwise stated, all data is taken from Fenner and

---

<sup>2</sup>However, in this case, the wheat “evolved” according to a controlled plan, rather than as a natural response to the evolution of the fly.



Ratcliffe (1965) and May and Anderson (1983).

### 5.2.1.1 Myxomatosis: An Illustrative Example

For the purpose of sport hunting, Thomas Austin introduced two dozen pairs of European rabbits (*Oryctolagus cuniculus*) into Australia in 1859. Extremely prolific and having no natural predators, rabbit numbers quickly exploded; reaching an estimated 300 million by the mid-twentieth century. Plaguing the entire continent, rabbits devastated Australian ecology and agriculture.

In a bid to control rabbit populations, the myxoma virus was deliberately introduced into Australia from South America in 1950. Exceedingly virulent, rabbit populations were decimated, with myxomatosis initially killing more than 99% of infected rabbits. However, in subsequent years, the severity of symptoms in rabbits and the proportion killed decreased (Lewontin, 1970). A similar pattern emerged when the same virus was introduced into France (1952) and Great Britain (1953).

To discover why rabbit deaths had diminished, an extensive series of laboratory tests were performed. Two primary factors emerged. Firstly, rabbit resistance had *increased*: injecting 1950 grade South American virus into 1957 Australian rabbits resulted in fewer deaths than 1950 Australian rabbits had suffered (Fenner & Ratcliffe, 1965). Secondly, the virulence of the virus had *decreased*: injecting 1957 Australian virus into 1957 South American rabbits resulted in fewer deaths than the 1950 virus had caused (Fenner & Ratcliffe, 1965). Whilst rabbit evolution towards increased resistance is unsurprising, to understand the less intuitive decrease in myxoma virulence, one must consider the strong coupling between the parasite's virulence and transmission.

Evidence suggests that high myxoma virulence is typically associated with an abundance of open lesions in infected rabbits. Such lesions enable transmission vectors—mosquitoes in Australia and France; fleas in Great Britain—to more easily bite infected wounds and acquire the virus. Low virulence is correspondingly associated with poor transmission. In general, parasites evolve to maximise their own transmission. Thus, by trading ease of transmission for a greater window of opportunity, the myxoma virus evolved towards the ESS level of moderate virulence.

It is understandable why this intermediate grade of virulence may be best: too high an  $\alpha$  [virulence] kills off hosts too fast, diminishing their capacity to transmit the infection; too low an  $\alpha$  is associated with a very quick recovery time, so that again transmission is relatively weak (May & Anderson, 1983).

Transmission vectors play a critical role in the evolution of virulence. Since mosquitoes

are more effective than fleas at reaching the blood supply of rabbits, fleas rely more heavily upon open sores. With higher virulence producing more open sores in infected rabbits, flea-transmitted myxoma has equilibrated at a higher level of virulence (i.e., a higher incidence of more virulent strains) in Great Britain than the mosquito-transmitted Australian and French strains.

Transmission by fleas is thought to be less efficient because the longer and larger mouthparts of mosquitoes enable the latter vector to transmit lower virulence grades of the virus more effectively (May & Anderson, 1983).

In biology, virulence evolves to the level that optimises parasite transmission. This varies between systems and over time. However, in artificial coevolution, the virulence of parasites is always high: parasites are encouraged to debilitate (beat in competition) a host *as much as possible*. Artificial analogues of the trade-off between “ease of transmission” and “window of opportunity” are not considered. As a result, artificial parasites are more likely to disengage from hosts—the equivalent of (sustained) high virulence myxomatosis driving rabbits extinct. In the event, both populations suffer.

### 5.2.1.2 A Case of Group Selectionism?

Several authors have attributed the myxoma virus’ decrease of virulence to group selection (e.g., Lewontin, 1970; Slatkin, 1983). Claiming that reduced virulence cannot be explained by individual selection alone, Lewontin (1970) argues the position that group selection *must* be invoked. Following the reasoning that each rabbit is a deme from the standpoint of the virus, he suggests:

... there is a tremendously high rate of deme extinction, with the result that those demes are left extant that are least virulent. This causes a general trend toward avirulence of the pathogen despite the complete lack of selective advantage of avirulence within demes (Lewontin, 1970).

Though highly unfashionable (refer to section 1.3.5.2), attributing the evolution of reduced myxoma virulence to group selection is not entirely unjustifiable.

Numerous theoretical studies have shown that group selection is theoretically possible when populations are extensively subdivided and local extinction is reasonably common [a dynamic particularly observed in parasites] (Slatkin, 1983).

However, whilst group selection is theoretically plausible, Lewontin's (1970) argument remains unfounded. Firstly, the mathematical models of May and Anderson (1983) demonstrate that individual selection *can* be exclusively used to explain the reduction in virulence. Secondly, the “general trend toward avirulence” is a false assumption. Whilst it is true that myxoma virulence fell, far from becoming avirulent, virulence stabilised at an intermediate value—even today, less than 10% of the UK's wild rabbits fully recover from myxomatosis. Finally, if each virus is considered as a deme in its own right (deme size tends to one), one should still see a reduction in virulence. Thus, since the (co)evolutionary process is unaffected by groups, by the “principle of parsimony” group selection should not be invoked.

Encouraging artificial parasites to sacrifice short-term advantage over hosts for the long-term good of both species, the arguments presented in this chapter (section 5.1) appear nakedly group selectionist. However, whilst this is true to some extent—“reducing the virulence” of artificial parasites *is* designed to encourage coevolutionary engagement—one can also consider the “virulence” of artificial parasites to be analogous to the transmissibility and infectiousness of biological parasites; but with the intricate couplings and complexities abstracted away.

Fortunately, when coevolutionary algorithms are applied to problem-solving, charges of group selectionism levied against “reduced parasite virulence” are largely irrelevant. The issue only becomes important if coevolutionary algorithms are used to simulate real host-parasite systems. In such cases, by directly adding parasite infectiousness and transmissibility to the model, “parasite virulence” should emerge.

### 5.2.2 The Laffer Curve

Describing the counterintuitive potential of government to *increase* revenue by *reducing* taxes, the infamous “Laffer curve” (figure 5.1) of macro economic theory exemplifies how short-term advantage can be sacrificed for long-term gain in socio-economic (people-state coevolutionary) systems.

Developed in the 1970s by Arthur Laffer, a Professor of business economics at the University of Southern California, the “Laffer curve” (as subsequently denominated by Jude Wanniski, then associate editor of the Wall Street Journal; Wanniski, 1978b) ushered in the “supply-side economics” that would provide the foundation for the Reagan-Thatcher revolution of the 1980s. At its core, the “Laffer curve” challenged Keynesian fiscal policy by suggesting that (under some circumstances) an increase in taxation may reduce government revenue. Making extensive usage of Wanniski (1978a), the principle




Figure 5.1: The “Laffer curve” showing the relationship between tax rate and revenue. At points *A* and *B*, the electorate desires more goods and services and is thus willing to pay more tax. At points *C* and *D* the electorate desires more private goods and services and thus wishes to pay lower tax rates. At point *E*, the “optimum” tax rate  $T^*$  returns the maximum tax revenue. Varying with the desires of the electorate,  $T^*$  may approach 100% during times of war (adapted from Wanniski, 1978a).

logic of the “Laffer curve” is described below.

Laffer noted that in any economy, ‘there are always two tax rates that yield the same revenues’ (Wanniski, 1978a). When tax rate is 100%, all production ceases in the monetary economy; thus, government revenue is 0. Existing largely to escape taxation, however, the barter economy will thrive (point *D* on the curve; figure 5.1). Conversely, if there is a 0% tax rate, there can be no government (since revenue is again 0). A state of anarchy exists (point *A*). At all taxation levels between these extremes, revenue is non-zero.

In between lies a curve. If taxation is reduced to a level less than 100% (point *C*, say), some level of the barter economy will be able to gain from the efficiencies of being in the monetary economy, despite the near-confiscatory tax rates. At the bottom end of the curve, if the people decide that some level of government is needed, a marginal tax rate will be introduced (point *B*, say). However, with the small loss of income from tax outweighing the efficiencies of being in the monetary economy, some segment of the economy will shift to barter. From the perspective of government, the “optimal” value of taxation—that which acquires the greatest revenue—is  $T^*$  (figure 5.1). Above this value taxation should be reduced; below this value it should be increased.

The optimal value of taxation,  $T^*$ , can be thought analogous to the virulence of biological parasites. However, rather than arising from the need for transmissibility and

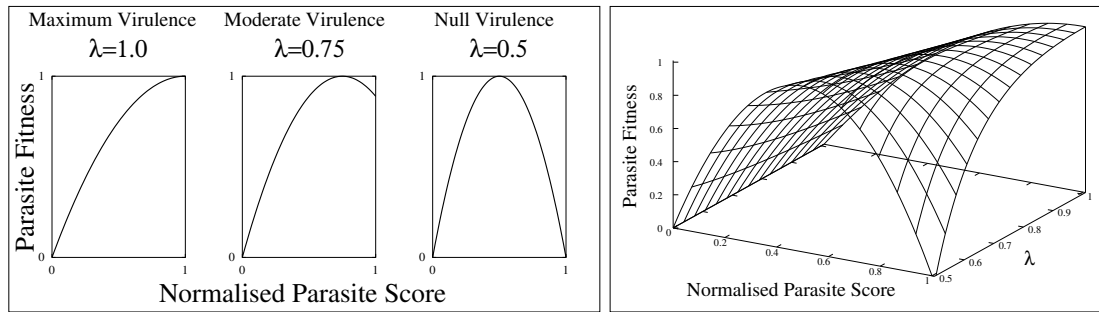


Figure 5.2: Parasite fitness,  $f_P$ , is a function of both parasite score,  $x$ , and virulence parameter,  $\lambda$ , (see equation 5.1). For convenience, the labels *maximum*, *moderate* and *null* virulence have been used for  $\lambda$  values 1.0, 0.75 and 0.5 respectively (left); however, there exists a continuum of possible virulence curves, producing a surface in three dimensions (right).

infectiousness,  $T^*$  results from the trade off between peoples’ need for security versus self-gain. Like virulence,  $T^*$  varies through time and between economies.

During times of war, for instance,  $T^*$  can approach 100%. Through the siege of Leningrad (8th September 1941 until 27th January 1944) the city produced for 900 days at tax rates of nearly 100%; soldiers and civilians worked to their limits in return for the basest of rations. Had citizens not wished to be taxed at high rate, the city would have fallen to Nazi Germany (Wanniski, 1978a).

When war stops,  $T^*$  falls rapidly; mirroring the electorate’s reduced demand for military goods and services. Following World War I, for example, having left wartime tax rates in place, the USA fell into recession in 1919-20. Running for President under a slogan promising a “return to normalcy”, William G. Harding was elected in a landslide victory. The subsequent tax rate reductions ushered in the economic expansion of the “roaring twenties” (Wanniski, 1978a).

In the following section, the analogies of biological virulence and the “Laffer Curve” are used to implement “reduced parasite virulence”; a novel technique designed to promote arms race engagement in artificial coevolutionary systems.

### 5.3 Implementation

Throughout this thesis, the term *score*—often calculated as the total number, or proportion, of victories in competition—is used to refer to the ability of an artificially coevolving individual to defeat the competitor(s) it is pitted against. Canonically, individuals receive fitness propensity,  $f_P$ , in proportion to their score. Here, however, this relationship is changed: to encourage coevolutionary engagement, parasites are encouraged to defeat

only  $\lambda * 100\%$  of host opponents; additional victories are *punished*.<sup>3</sup>

Parasite scores are normalised with respect to the maximum score achieved that generation such that the best current parasite *always* achieves a score of 1. Parasite fitness,  $f_P$ , is then calculated as a function of score,  $x$ , and virulence,  $\lambda$  ( $0.5 \leq \lambda \leq 1.0$ ), such that:

$$f_P(x, \lambda) = \frac{2x}{\lambda} - \frac{x^2}{\lambda^2} \quad (5.1)$$

Thus, a parasite achieves optimum fitness by winning a proportion of contests equal to a fraction,  $\lambda$ , of that achieved by the best parasite. By varying  $\lambda$ , parasites can be encouraged to be more, or less, virulent (see figure 5.2). Although there exists a continuum of possible curves, throughout this thesis three values of  $\lambda$  are used most often. These are labelled as follows: *Maximum* virulence ( $\lambda = 1.0$ , the equivalent of canonical parasites) where parasites are encouraged to beat as many hosts as possible; *Moderate* virulence ( $\lambda = 0.75$ ) where parasites are encouraged to achieve a win-rate three-quarters that of the highest scoring current parasite; and *Null* virulence ( $\lambda = 0.5$ ), where the fittest parasites achieve half the win-rate of their highest-scoring conspecifics. Notice that a value of lambda less than 0.5 would encourage *cooperation* between parasites and hosts; with parasites continually striving to achieve more losses than wins.

It is important to emphasise at this point that drawing inspiration from natural systems does not make reduced virulence more theoretically grounded than alternative approaches (such as the phantom parasite or  $\Phi$  function; section 3.4.5, figure 3.4); nor does it imply that biological systems necessarily optimise. As an alternative analogy, reducing virulence can be thought of as maintaining a gradient for selection; forcing parasites to evolve in difficulty at a similar speed to hosts. For example, in order to maximise pupils' learning, teachers must teach material that is neither too difficult nor too easy. Rather, teaching material should be presented on a smooth, continuous gradient of difficulty, consistent with the current academic requirements of pupils (e.g., Sklar & Pollack, 2000).

### 5.3.1 Maintaining Relative Fitness Diversity

Since disengagement occurs when individuals can no longer be discriminated on the basis of fitness alone, an immediate question arises: in order to produce discriminatory parasites, why not encourage them to maximise fitness variation in hosts? Rather than utilise reduced virulence, why not directly equate parasite fitness with the standard deviation of opponent hosts' scores? This section suggests an answer.

---

<sup>3</sup>“Reduced virulence” may be applied to *either* coevolving population (or both). However, for brevity, assume here that it affects only one population: the “parasite” population.

Firstly, reduced virulence and direct fitness diversity encouragement may often behave similarly. For instance, assume parasite  $P$  records “win, draw, loss” against three hosts; then (assuming  $\lambda = 0.5$  and *at least* one parasite wins all competitions),  $P$  is optimal under both schemes. Yet, whilst both schemes similarly encourage fitness diversity, whether implicitly or explicitly, reduced virulence is (usually) less computationally expensive to implement.

More importantly,  $\lambda$  enables reduced virulence to be “tuned”, as desired, to fit the requirements of each particular coevolutionary system.<sup>4</sup> It is not immediately obvious how direct diversity encouragement could effectively incorporate such fine control. As a result, reduced virulence offers a major advantage over fitness diversity encouragement.

### 5.3.2 Asymmetry

Often, coevolutionary systems are asymmetric: populations may differ genetically (in terms of encoding) or behaviourally (in terms of goal strategy). Such asymmetry may result in an inherent advantage for one population. When coevolving pursuers and evaders, for example, it is often much easier, at least initially, to be a successful evader (Cliff & Miller, 1995). Given that disengagement results from one population out-performing the other, it is intuitive that an inherent asymmetrical advantage toward a particular population will encourage the likelihood of coevolutionary disengagement.

However, if parasite virulence were reduced, this trend could be reversed. By rewarding parasites capable of discriminating hosts—those that occasionally lose—with *more* offspring, asymmetrical advantage may be reduced; thus reducing the likelihood of disengagement. As such, the optimum level of virulence for each system will vary with the nature of asymmetry: large asymmetries may require low virulence, and *vice versa*. Critically, preventing disengagement will improve coevolutionary optimisation if a reduction in periods of degrading coevolutionary drift can be achieved without sacrificing the selection pressure that ensures progress.

## 5.4 Summary

Inspired by the virulence of natural parasites, “reduced parasite virulence”, a novel technique designed to prevent disengagement, has been presented.

---

<sup>4</sup> $\lambda$  has the potential for self-adaptation (refer to section 6.4) and may even be evolved.

# Chapter 6

## Combating Disengagement

---

Disengagement (introduced in detail in section 3.4.5) is a detrimental phenomenon that hinders the progress of coevolutionary systems. Occurring when individuals are unable to be discriminated on the basis of fitness alone—usually due to one population outperforming the other—disengagement removes the selection pressures on each population, producing (random) coevolutionary drift. Often, such drift is derogatory; particularly when disengagement occurs once populations contain above average solutions.

Inspired by the behaviour of natural host-parasite systems, the previous chapter introduced *reduced parasite virulence*, a technique specifically designed to combat coevolutionary disengagement. By reining in the inherent advantage one population (the parasites) may hold, reduced virulence attempts to stop the out performance of one population (the hosts) by another (the parasites). If successful, within-population fitness diversity is maintained and selection pressures persist: coevolutionary progress results.

In this chapter, the ability of reduced virulence to combat coevolutionary disengagement is analysed through a series of experiments in a number of problem domains. The reader should be aware that the aim of these experiments is to tease out the general effects reduced virulence has on the underlying dynamical behaviour of each system: there is no attempt to find novel solutions to specific optimisation problems; particular solutions are unimportant in themselves. However, it is hoped that any conclusions drawn will be generalisable to other domains.



## 6.1 The Numbers Game

### 6.1.1 Background

First introduced by Watson and Pollack (2001), the *numbers game* (also referred to as the “counting ones problem”; e.g., Cartlidge & Bullock, 2002, 2003) was initially developed to demonstrate the potential for minimal evolutionary substrates to exhibit the detrimental coevolutionary dynamics of disengagement, overspecialisation and cycling.<sup>1</sup> By stripping away “unnecessary” domain complexity, Watson and Pollack (2001) produced a valuable coevolutionary substrate that is both simple to visualize and analyse, yet capable of exhibiting a range of complex dynamics. As a result, the numbers game has quickly become established as a standard coevolutionary modelling domain (adopted by Cartlidge & Bullock, 2002; Bucci & Pollack, 2003; Ficici & Pollack, 2003; De Jong & Pollack, 2004).

The numbers game, as implemented by Watson and Pollack (2001), has three alternative formats, each designed to highlight a different coevolutionary dynamic. However, for disengagement, only the simplest 1-dimensional game is required. Here, the goal is to maximise an evolving integer value,  $n$ . Each generation, individuals receive a relative fitness score based upon the proportion of opponents (randomly sampled) they are able to beat: victories are awarded if an individual’s  $n$  value is *greater than* that of the opponent. With individuals each encoded as a 100 bit binary string, the  $n$  value of each is calculated as the total number of 1s in the string. Thus, the maximum value of  $n$  is 100. Surprisingly, whilst this integer maximisation problem is trivial for standard evolutionary approaches—consider an objective fitness function  $f(x) = n$ , for example—coevolutionary systems can fail. This allows interesting system dynamics.

In this section, an adaptation of the numbers game is utilised to demonstrate the effect that reducing virulence has upon disengagement when there exists an asymmetrical advantage favouring one coevolving population. As discussed in chapter 5, it is often the case that one side of a coevolutionary contest has an (at least temporary) advantage over the other in terms of the ease with which successful counter-adaptations are discovered. To create asymmetry, one population (the parasites) is given a favourable mutation bias such that each point mutation is more likely to produce a 1 than a 0; the other population (the hosts) mutates to either bit with *equal* likelihood. As such, it is *easier* for parasites to evolve higher values of  $n$ .

In complex coevolutionary systems asymmetrical advantage is free to fluctuate in am-

---

<sup>1</sup>To the author’s knowledge, this was the first successful visualization of disengagement to appear in the coevolutionary literature.

plitude and shift between populations. However, in order to aid analysis, the asymmetry introduced into the numbers game, in this section at least, is fixed throughout each run. Although fixed asymmetry may not be representative of coevolutionary asymmetry in general, it affords greater simulation control whilst behaving in a manner similar (for the purposes of this thesis) to more realistic asymmetry.

## 6.1.2 Investigating Engagement

### 6.1.2.1 Set Up

Two reproductively isolated populations of size 25 are coevolved. Individuals in each population consist of binary strings containing 100 bits, with each bit initialised to 0 in generation 0. As discussed above, the goal of the numbers problem is to evolve strings containing as many 1s as possible. Of course, in this *toy* domain, as observers we can assess the absolute fitness or objective quality of each individual by counting its 1-alleles. This affords a useful way of measuring progress. However, it should be noticed that the coevolutionary algorithm does not make use of this absolute measure, only having access to the relative fitness measure described below.

Members of one population are selected to play a set of *pair-wise* contests against a random sample of 5 opponents from the competing population. During each contest, the individual with the genotype containing the greatest number of 1-alleles receives a fitness point. Each opponent receives half a fitness point for contests resulting in a draw. Individuals in both populations reproduce asexually with parents chosen through tournament selection (tournament size 5; winner reproduces). Offspring have a small probability of mutation,  $m$ . Unless specified otherwise, the probability of mutation at each locus,  $m$ , is 0.03.

An asymmetry was introduced by varying mutation bias  $B_{par}$  ( $0 \leq B_{par} \leq 1$ ) in favour of one population, henceforth classified as the *parasites*. Given mutation at a particular parasite locus, the substitution of a 1 or 0 occurs with probability  $B_{par}$  and  $1 - B_{par}$ , respectively. In contrast, the coevolving *host* population substitutes a 0 or 1 with *equal likelihood* whenever mutation occurs. We thus see that if  $B_{par} > 0.5$ , there is a bias in favour of evolving parasites with more ones—an asymmetry that favours the parasite population.

Two parasite mutation bias values and two parasite virulence levels were tested over a series of runs;  $B_{par} = 0.75, 0.9$  and  $\lambda = 0.75, 1.0$ : Moderate and Maximum virulence, respectively (refer to section 5.3). Unless otherwise stated, the value of  $\lambda$  remained constant throughout each run.




Figure 6.1: Typical numbers game coevolution ( $B_{par} = 0.75$  and  $m = 0.03$ ), using the same initial conditions each run. With Maximum virulence parasites (left) there are two periods of disengagement. The second period may be prevented by switching to Moderate virulence at generation 250 (centre). Populations remain engaged throughout the entire run when Moderate virulence is utilised from the beginning (right).

### 6.1.2.2 Results

Figures 6.1 and 6.2 each display three typical runs, using a parasite mutation bias,  $B_{par}$ , of 0.75 and 0.9 respectively. The large graphs display the  $n$  value of each individual each generation—parasites are shown light grey, hosts are dark—allowing progress to be easily monitored. The two smaller graphs, below, display the mean relative fitness of each population: the average number of victories individuals in each population secure each generation. Whilst these graphs afford less insight into coevolutionary progress, they accurately reflect the information available to the coevolutionary algorithm, and present an illustrative example of the inherent difficulties of using standard (relative) “fitness over time” graphs to visualize coevolutionary progress (refer to section 3.4.1 for a detailed discussion). This style of graph (inherited directly from Watson & Pollack, 2001) will be used several times throughout this chapter.

When employing maximally virulent parasites (figure 6.1, left) the populations have a tendency to disengage. This can be observed between generations 150 – 175 and again between 250 – 500. During these periods of disengagement the populations drift back to their relative baseline performance, equal to the mutation bias,  $B_{par} = 0.75$  and  $B_{host} = 0.5$ . Only once the populations re-engage by chance is there an improvement in absolute fitness. Repeating the run with the same random seed, the second period of disengagement depicted in figure 6.1 is prevented if Moderate virulence ( $\lambda = 0.75$ ) is introduced at generation 250 (figure 6.1, centre). Notice that the left and centre graphs are identical until generation 250—the point at which parasite virulence is changed to Moderate. In contrast to Maximum virulence, when Moderate parasites are used from the start (figure 6.1, right),

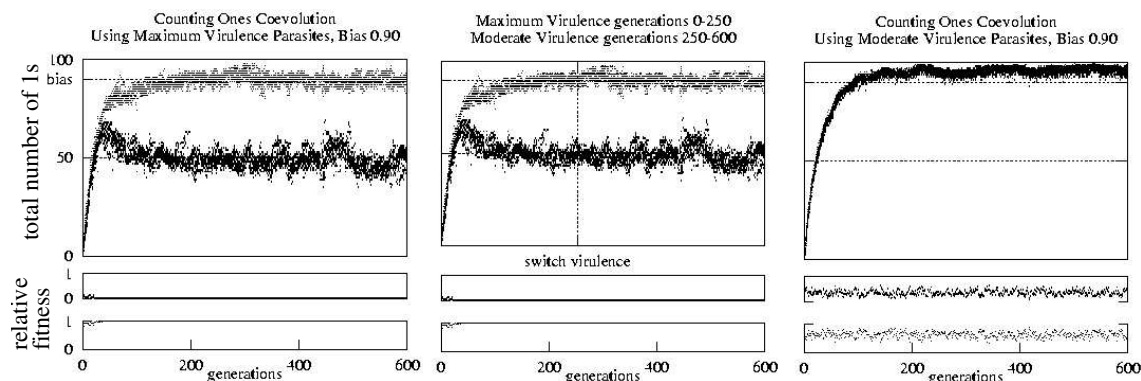


Figure 6.2: Typical numbers game coevolution ( $B_{par} = 0.90$  and  $m = 0.03$ ), using the same initial conditions each run. With Maximum virulence parasites (left) the populations disengage within 50 generations and fail to re-engage. Switching to Moderate virulence during disengagement (centre) has no effect. Utilising Moderate virulence from the beginning of a run (i.e., *before the occurrence* of disengagement) enables the populations to remain engaged throughout the run (right).

the populations remain engaged throughout the entire run, achieving a continuously high level of performance.

With a parasite mutation bias of 0.9, the increased asymmetry exacerbates the effects of reducing virulence. Typical of all runs, figure 6.2, left, shows that a bias of 0.9 is too great for the host population to remain engaged with maximally virulent parasites after the initial 50 generations. However, runs employing Moderate virulence maintain population engagement despite the underlying asymmetry (figure 6.2, right). It should be noted, however, that whilst Moderate virulence helps to prevent disengagement from occurring, it does not encourage populations to re-engage. Switching to Moderate virulence during coevolutionary disengagement has no affect (figure 6.2, centre; see also section 6.1.3.3).

These results are sensitive to variation in both population size and the number of opponents played by each individual. As either parameter increases, the probability of disengagement decreases due to the increased frequency of meeting varied opponents. However, there remains a chance of disengagement even when sample and population sizes become very large—the phenomenon does not disappear. The results observed in this section are qualitatively robust to mutation rate,  $m = [0.005, 0.05]$ , and tournament size,  $T = [2, 15]$ .

The results displayed in figures 6.1 and 6.2 clearly demonstrate that reducing parasite virulence in asymmetric coevolution can reduce the effects of disengagement. In particular, the greater the inherent asymmetry, the greater the effect reducing virulence has upon results. The asymmetry imposed in this model gave the biased parasite population a great advantage over the coevolving host population. Purely by stochastic effects one

would expect individuals from parasite populations to contain more 1s than those from host populations. This is observed in figures 6.1 and 6.2. The difference in the speed with which the two populations initially move through the genotype space (resulting from the different mutation biases) ensures that disengagement occurs rapidly, and once it has occurred the same mutation biases tend to restrict each population to a different portion of the genotype space. Mutation bias pushes each population towards a particular ratio of 1s to 0s, i.e., 0.5 for hosts and, dependent upon mutation bias, 0.75 or 0.9 for the parasites. As a result, both populations will remain disengaged until the gap between them is, at least temporarily, bridged by the occurrence of a very large number of mutation events, e.g., figure 6.1, left, generation 500.

The first generation of parasite offspring will on average contain many more 1s than that of the host population. However, under reduced virulence, any parasite that beats all opponents is *less* fit than those parasites that lose a small percentage of contests. In this way, acceleration is decreased as the parasites resist their mutation bias. Moderate virulence parasites appear to actively prevent disengagement. Using the continued selection pressure ensured through engagement, hosts evolve to a higher objective quality than would otherwise be possible. It should not be overlooked, however, that Moderate parasites gain from this relationship too, as both populations evolve to a greater standard than either would alone (figure 6.2, right). However, as parasite virulence is decreased there is a tendency for coevolution to stagnate at a sub-optimal but highly engaged fluid local optimum. In order to push populations to optimal solutions, stronger selection pressure is required (see section 6.1.3.2, below).

### 6.1.3 Further Investigations

In the previous section, the ability of reduced virulence to combat disengagement in the numbers game was tested. In this section, further investigations are performed in order to uncover exactly when reduced virulence is beneficial, how the value of  $\lambda$  should be chosen and how a reduction in virulence compares to alternative engagement-maintenance algorithms. Throughout this section, unless otherwise stated, a mutation bias for parasites  $B_{par} = 0.75$ , and per locus mutation rate  $m = 0.03$  is used. Once again, unless otherwise stated, the sample size of opponents is 5 and tournament size for selection  $T = 5$  (winner always chosen to reproduce).

### 6.1.3.1 Comparing the Phantom Parasite

Of the engagement-maintenance techniques discussed in section 3.4.5, the only domain-independent algorithm versatile enough to fit the two-population coevolutionary model used for the numbers game is the Phantom Parasite (Rosin, 1997, figure 3.4, left). This section compares reduced virulence with the Phantom Parasite to observe the behaviour of both algorithms when subject to varying levels of asymmetry, mutation and selection pressure. Maximum virulence ( $\lambda = 1.0$ ), Moderate virulence ( $\lambda = 0.75$ ), Null virulence ( $\lambda = 0.5$ ) and the Phantom Parasite were each tested over a series of 50 runs; with parameter values  $B_{par} = [0.5, 0.99]$ ,  $m = [0.005, 0.050]$ , and  $T = [5, 15]$ . Figure 6.3 displays the *mean* number of disengaged generations occurring each run.

Each graph clearly demonstrates the relationship between asymmetry and disengagement; irrespective of the engagement-maintenance technique, greater asymmetry (parasite bias) produces more disengagement. In contrast, disengagement occurs less frequently as virulence is reduced. Thus, reducing virulence is particularly beneficial when asymmetry is high. As the mutation rate increases, the effects of asymmetry are exacerbated: given that the inherent asymmetrical advantage in favour of parasites is a mutational bias, any increase in mutation rate directly increases asymmetry. Similarly, an increase in selection pressure in favour of parasites ( $T_p - T_h$ ) also exacerbates the effects of asymmetry, since “weaker” individuals are less likely to breed.

As one would expect, the Phantom Parasite performs better than canonical parasites ( $\lambda = 1.0$ ) across all levels of asymmetry, mutation and selection pressure. However, it can be observed that the Phantom Parasite is *sensitive* to mutation rate—rather than exhibiting a monotonic relationship between disengagement and mutation rate, disengagement increases as  $m$  diverges from 0.020. The Phantom Parasite thus behaves less predictably than reduced virulence, perhaps due to the discontinuity in the gradient of the fitness curve (see figure 3.4, left). In contrast, the Phantom Parasite is *not* as effective as Moderate or Null virulence in reducing disengagement when asymmetry is high. Again, this is as we would expect. The Phantom Parasite is *most like* reduced virulence with  $\lambda = \frac{N-1}{N}$  (thus sensitive to sample size  $N$ , here  $N = 5$ ); with the exception that it is *always* best to win all competitions. As such, the Phantom Parasite, here similar to  $\lambda = 0.80$ , performs better than Maximum but worse than Moderate and Null virulence.

Thus, in comparison with the Phantom Parasite, reduced virulence acts more predictably (is less sensitive to coevolutionary parameters), produces less disengagement (assuming  $\lambda$  has been chosen adequately) and is more flexible (i.e.,  $\lambda$  can be varied as required).

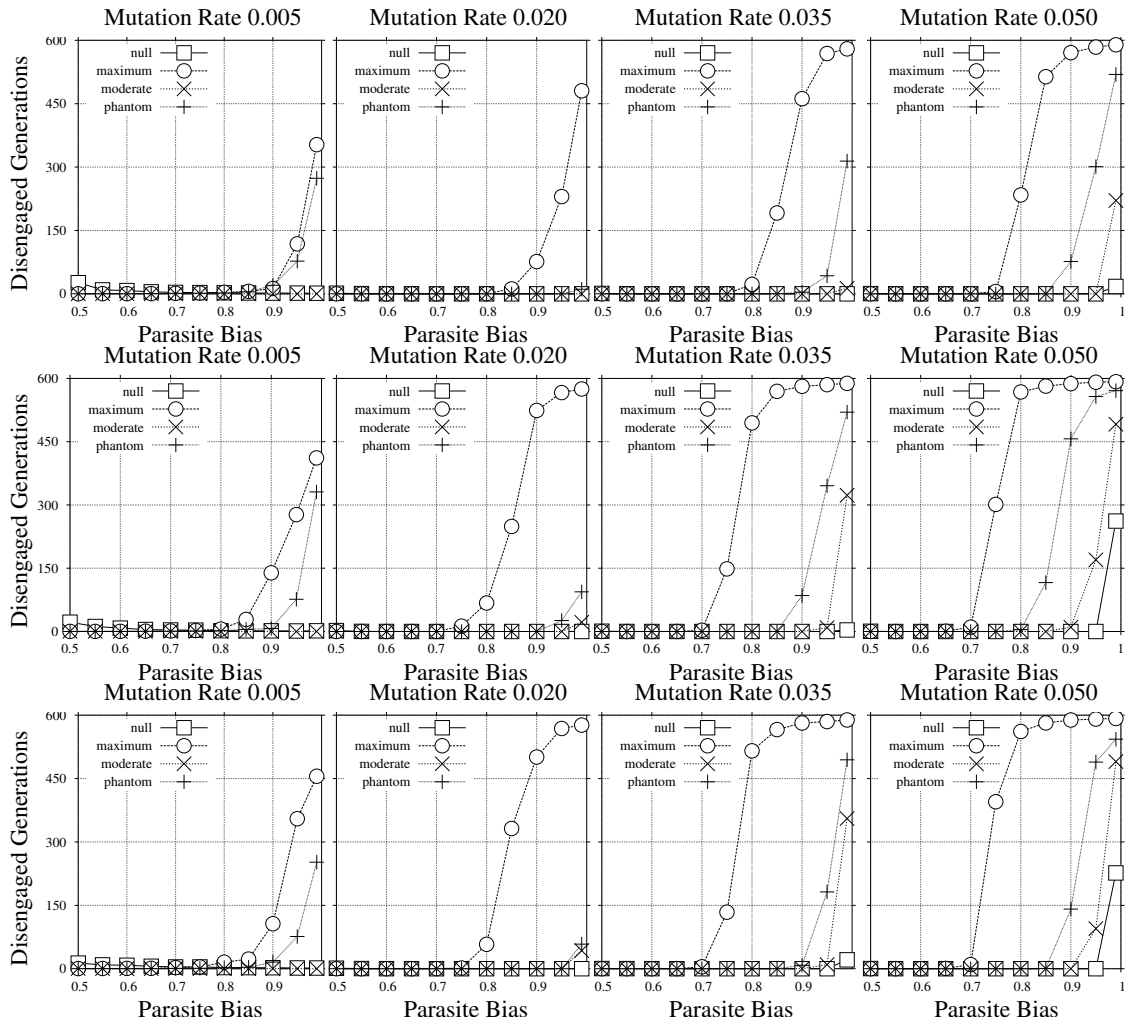


Figure 6.3: The average of 50 runs (each 600 generations) of the numbers game, with mutation rate ranging between  $m = 0.005$  (far left) and  $m = 0.050$  (far right). From top to bottom, selection pressure is increased in favour of parasites (top: host tournament size  $T_h = 15$ , parasite  $T_p = 5$ ; middle:  $T_h = 5$ ,  $T_p = 5$ ; bottom:  $T_h = 5$ ,  $T_p = 15$ ). As expected, an increase in  $B_{par}$ ,  $m$ , or  $T_p - T_h$ , increases the likelihood of disengagement. In general, reducing virulence succeeds in reducing the number of disengaged generations: in particular the greater the asymmetry, the greater virulence should be reduced. The phantom parasite performs less well and is clearly sensitive to mutation rate.

### 6.1.3.2 The Engagement-Optimality Trade Off

The reduced virulence technique has been observed to limit the effects of disengagement, however, the effect upon performance (in this case, the total number of 1-alleles within the genome of the best host) has not been considered. As reducing virulence directly *interferes* with the selection pressure upon parasites, it is likely that the performance of the coevolutionary system will be affected in some way.

Figure 6.1 hints at an answer. When utilising Maximum virulence parasites (figure 6.1, left), the system briefly reaches near optimal performance (generations 100-150) before disengaging. However, when using Moderate virulence parasites (centre and right), maximum performance is sub-optimal, despite being much more stable. Since Maximum virulence parasites result in the strongest selection pressure upon hosts, the following hypothesis is proposed: given that a coevolutionary system can maintain engagement for a sufficient period of time, Maximum virulence will, on average, push a system to a greater level of performance than can be achieved by parasites with reduced virulence.

This hypothesis was tested by *forcing* populations to remain engaged throughout a run. In a variant of the numbers game, each member of a single population is assessed against a random sample of opponents drawn from the *same* population. In this way, to the extent that the population remains phenotypically diverse, it must remain engaged with itself over evolutionary time.<sup>2</sup>

For each virulence level, 30 runs were performed (population size 25). After approximately 100 generations the coevolutionary system settled into equilibrium at its highest level of performance. With a 95% confidence limit, at generation 600, the population mean level of performance for each virulence scheme was: Maximum  $99.1 \pm 0.38\%$ , Moderate  $94.8 \pm 1.50\%$  and Null  $75.1 \pm 2.59\%$ . It appears that, if engagement can be guaranteed, maximum performance increases with virulence.

Thus, performance is maximised by setting virulence to be as great as possible, without sacrificing engagement. The trade off between engagement and performance must be balanced by choosing the *optimal* level of virulence,  $\lambda$ . However, once a system has disengaged, it is not obvious which virulence level is optimal. This is discussed in the following section.

---

<sup>2</sup>This is true since the numbers game, as stated, is 1-dimensional. With multiple dimensions, it is possible for certain dimensions to disengage; perhaps leading to overspecialisation (section 3.4.2).



### 6.1.3.3 Re-engagement

The results in section 6.1.3.1 show that reducing virulence can limit the propensity for asymmetric populations to disengage. However, as yet, no insight has been gained into the effects that parasite virulence has upon populations that are *already* disengaged. Is it possible that certain virulence functions increase population diversity and so encourage populations to re-engage? To test this hypothesis, populations were forced into disengagement at the start of each run. This was achieved by initialising every individual in the host population to a genotype consisting of 0s, whilst initialising individuals in the parasite population to have some integer,  $n > 0$ , number of 1s in their genomes. In each run, the number of generations until first engagement was recorded across a range of values of  $n$ . Using the same initial conditions (i.e., same random seed), 30 runs were performed using Maximum, Moderate and Null virulence parasites. Results demonstrated absolutely no difference in the time to re-engagement—the time until populations first engage is independent of virulence: it is purely a stochastic process. The host and parasite populations drift until an instance arises when at least one parasite does not score maximally against at least one host. This is a re-engagement event.

Although the time to first engagement is independent of parasite virulence, the behaviour of the system henceforth most certainly is not. Imagine a situation where every parasite but one scores maximally and every host but one scores minimally. With a Moderate or Null virulence scheme the non-maximal scoring parasite would be rewarded with greater reproductive success. The same would be true of the non-minimal scoring host. Effectively, the populations would be drawn together and engagement would be encouraged. With a Maximum virulence scheme, however, the situation would be very different. Rather than be rewarded with progeny, the non-maximum scoring parasite would be unlikely to have any progeny at all whilst every other parasite would produce marginally more offspring than before. Effectively, we would see the populations *bounce apart*—disengagement is likely to re-occur the very next generation.

### 6.1.3.4 Random Parasites: Canonical Evolution

This section considers how coevolution compares to canonical evolution in the numbers game. Canonical evolution is implemented by evolving a population of solutions (hosts) against non-hereditary test-problems (parasites), randomly selected at the beginning of each generation with a uniform distribution of densities (total number of 1s in the genome).

Figure 6.4 displays the output of one typical run. As can be seen, the host population




Figure 6.4: Random parasites with no heredity. Although most parasites are either “upper” or “lower” parasites, the likelihood is that some parasites will fall within the set of strategies capable of *discriminating* between hosts. This produces a small selection pressure encouraging the objective quality of the host population to fluctuate around 10% higher than that achieved by mutation bias alone.

fluctuates in performance around 60% 1-alleles—10% higher than the population would tend to by mutation alone—implying Random parasites are exerting some selection pressure. Consider those parasites that fall *above* the host population: the *upper* parasites. By definition, each competition between a host and upper parasite will result in victory for the parasite. Conversely, consider those parasites falling below the host population: the *lower* parasites. In this case the reverse will be true. Each competition will result in a win for the host. Thus, how strong or weak a host is in relation to other hosts cannot be discerned on the basis of competitions with either upper or lower parasites. The host population is disengaged from the upper and lower parasites.

It is likely, however, that some parasites will engage with the host population: *discriminating* parasites. These individuals are able to discern, to some degree, the relative strength of hosts, thus producing systematic selection pressure. However, due to the relatively tight distribution of hosts—in terms of *possible* distributions, the host population is relatively converged—only a small proportion of parasites are able to discriminate. As such, Random non-hereditary parasites exert a constant, but relatively weak, selection pressure—the system has a *small degree* of engagement.

The impetus behind coevolving test-problems with solutions derives from the desire to automatically sample tests at the *required* difficulty, reducing the necessary number

of tests and resulting in computational efficiency. The effectiveness of coevolution in improving performance can be observed by comparing figure 6.1, right, and figure 6.4: coevolving hosts fluctuate around 30% higher than those evolving with Random parasites.

However, the occurrence of disengagement drastically changes this relationship. Disengaged populations experience *no* selection pressure and thus drift to their respective base-line levels of performance. Fluctuating at around 50% 1-alleles, disengaged hosts perform *worse* than hosts assessed through standard evolution. However wasteful Random parasite populations are in terms of discriminatory ability per parasite, the small but continuous selection pressure they exert is better than long periods with no selection pressure at all. Although coevolution may have the ability to outperform standard evolution in certain domains, long episodes of disengagement can easily reverse this advantage.

### 6.1.4 Conclusions

Several conclusions can be drawn from the numbers game experiments, however, one cannot assume that these are true *in general*. Firstly, there exists a trade off between reducing virulence to encourage engagement and increasing virulence to improve performance. Secondly, disengagement can reduce the performance of coevolutionary algorithms to below that of standard evolution, even if the domain is particularly *suited* to a coevolutionary approach. Once disengagement has occurred, the level of virulence is irrelevant; re-engagement occurs stochastically. *Nevertheless*, reducing virulence increases the likelihood of prolonging re-engagement once it occurs. Finally, reduced virulence outperforms its closest competitor, the Phantom Parasite.

## 6.2 The Coin Toss Game

### 6.2.1 Diversity Maintenance

Disengagement occurs when intra-population fitness diversity reduces to zero. Moderating virulence counter-acts disengagement by selecting, for reproduction, parasites that are occasionally beaten. This preserves a selection gradient for hosts which, in turn, maintains relative fitness diversity in both populations.

A tendency towards reduced *population* diversity (and the associated problem of premature convergence) has long been a major concern of the evolutionary computation research community. As such, a suite of diversity maintenance techniques have been proposed, including e.g., deterministic crowding (De Jong, 1975), explicit fitness sharing (Goldberg & Richardson, 1987), competitive fitness sharing (Rosin, 1997; Rosin &

Belew, 1997), resource sharing (Juillé & Pollack, 1998b), and spatial embedding (e.g., Hillis, 1990). These approaches are attempts to maintain genetic diversity on the assumption that a loss of diversity can be harmful to optimisation as it may restrict search to local optima (for a more comprehensive overview refer to section 3.4.2).

Resource (or competitive fitness) sharing maintains genetic diversity in a population by encouraging niching: individuals are rewarded for being able to solve tests that few others can. This idea has been extended to coevolutionary scenarios where opponents are treated as a commodity or resource. Rather than gain a fitness point for each victory against an opponent (simple fitness), one fitness point is shared among the competitors that beat a particular individual. Thus, individuals are rewarded less for *how many* opponents they beat and more for *who* they beat, rewarding phenotypic diversity and maintaining genetic diversity.

Since disengagement is associated with a loss of diversity, could it be prevented by simple diversity maintenance approaches? Perhaps reducing virulence is only preventing disengagement by mimicking these existing techniques? If so, it is largely superfluous. The following study contrasts reduced virulence with resource sharing in order to explore whether they are effectively the same, or different, in some fundamental sense.

### 6.2.2 Set Up

To compare the influence of parasite (and host) resource sharing with that of reduced parasite virulence, the numbers game does not offer enough complexity. Specifically, to develop niching, resource sharing requires multiple dimensions of play. Here, a simple (novel) evolutionary domain is presented: the *coin toss game*. Like the numbers game, the coin toss game is deliberately designed to be simple to analyse and visualize, whilst allowing potentially complex coevolutionary dynamics to occur; including cycling, over-specialisation, disengagement and mediocre stability. The coin toss game is antagonistic. Parasites each toss a coin; the result of which is probabilistically determined by the genetic structure of the parasite. Depending upon the result of the toss, a host opponent is rewarded for matching either Head-alleles, or Tail-alleles, of the parasite. If matched, a parasite is punished.

Games with this type of dynamic often suffer from coevolutionary cycling, as hosts chase parasites through the strategy space. Although desirable generalist strategies exist, populations are easily diverted from them as they exploit the temporary idiosyncrasies of their opponents. Resource sharing is one way of discouraging this type of short-termist behaviour. By maintaining a diverse strategy-base in each population, the value of ex-

exploiting idiosyncrasies is reduced; encouraging generalists. Unfortunately, an alternative mediocre stable scenario is possible in which populations “speciate” such that they exhibit a number of different sub-optimal strategies that together form a stable combination: collusive disengagement (refer to section 3.4.5.1). In this sense, the game is similar to any number of scenarios in which a generalist strategy is desirable (from the perspective of the observer), but difficult to evolve in practise—e.g., scissors-paper-stone, immune systems, etc.

Two distinct populations of size 50 are coevolved—hosts and parasites. Individuals in each population consist of binary strings containing 100 bits, initialised *randomly* in generation 0. Each generation, members of the host population are selected to play a set of pair-wise contests against a random sample of 10 opponents from the parasite population. The aim for hosts is to match as many parasite alleles as possible. Antagonistically, parasites aim to mismatch host alleles. Both populations breed asexually, with each individual having a small probability of unbiased mutation per locus,  $m = 0.03$ . Tournament selection is used (tournament size 5) with the winner of each tournament always chosen to reproduce.

Not all loci are involved in allele matching. For parasites with many H-alleles, allele matching tends to involve only those loci at which the parasite possesses H-alleles. For parasites with many T-alleles, the game tends to involve only those loci at which the parasite possesses T-alleles. Whether H-allele loci or T-allele loci are involved is determined probabilistically. The probability,  $p(\text{Head})$ , a parasite will toss a Head—thus forcing host opponents to match H-alleles—increases with the total number of H-alleles,  $x$ , in the parasite’s genotype, such that:

$$p(\text{Heads}) = \frac{1}{2} \left\{ 1 + \tanh \left( \frac{x - 50}{7} \right) \right\} \quad (6.1)$$

Once the coin has been tossed, a host wins by matching alleles in *at least*  $M = 30$  loci, else the parasite wins (see figure 6.5).

Having several antagonistic points of attraction, the coin toss game is designed to exhibit interesting coevolutionary dynamics. Mutation bias attracts both populations towards genotypes containing 50% H-alleles and 50% T-alleles. However, given a parasite tosses Heads (or Tails), it is advantageous for the host to have *as many* H- (or T-) alleles as possible. Thus, the host population is attracted towards homogeneous genotypes (all Hs or all Ts). The direction of attraction for hosts (towards either 100% H- or T-alleles) depends upon the frequency with which the parasite population tossed either Heads or Tails. This occurs with increasing frequency the further parasite genotypes vary from a 50% H/T mix. Thus, parasites are also attracted away from a 50% mix, but in the *opposite*

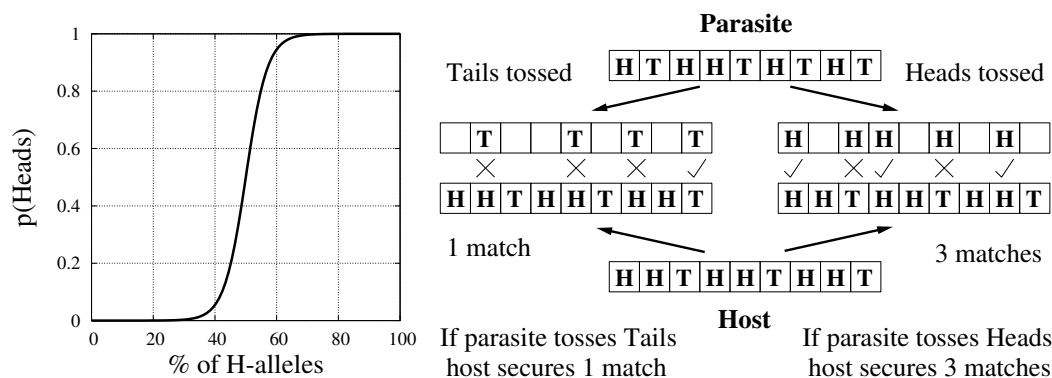


Figure 6.5: The probability  $p(\text{Heads})$  of parasites tossing Heads (left, see also equation 6.1). Depending upon whether parasites toss Tails or Heads, hosts attempt to match *either* T- or H-alleles (right). If  $\text{matches} \geq M$  then the host wins, otherwise the parasite wins.

direction to hosts. Parasites deviating too far from 50%, however, become *too predictable*. In general, the most difficult parasites to match are those having approximately 50% H- and 50% T-alleles.

This matching game resembles the density classification task for 1-D cellular automata, for which Juillé and Pollack (1998a, 1998b) and Pagie and Hogeweg (2000) utilised a method of virulence reduction—the  $\Phi$  function (refer to section 3.4.5). The density classification task for cellular automata is difficult—no rule set exists which can correctly classify all ICs (Land & Belew, 1995)—as such, consistently coevolving two populations towards continuous improvement is problematic (Paredis, 1997). Whilst coevolving CA rules, Juillé and Pollack found it necessary to “reduce the virulence” of ICs in order to stop disengagement; despite the use of resource sharing. Here, the contribution of two approaches to improving coevolutionary optimisation—resource sharing and reduced parasite virulence—are teased apart.

Two  $\lambda$  values were tested over a series of runs;  $\lambda = 1.0$  (Maximum) and  $\lambda = 0.5$  (Null). The value of  $\lambda$  remained constant throughout each run. Runs were performed under four conditions: Maximum virulence without resource sharing (i.e., standard coevolution); Maximum virulence with resource sharing; Null virulence without resource sharing; both Null virulence and resource sharing. Under each condition, the degree of niching or genotypic diversity within each population was calculated using a *linkage disequilibrium* measure (e.g., Barton & Gale, 1993; Falconer, 1996) particularly sensitive to the effects of resource sharing.

### Linkage Disequilibrium:

Linkage disequilibrium is a quantitative genetics measure used to determine a popula-

tion's distance from equilibrium. Under normal conditions, a well mixed (sexual) population will tend towards (equilibrate at) random combinations of genes, in Hardy-Weinberg proportions. Disequilibrium occurs when traits are strongly associated, perhaps arising from intermixture of populations with different gene frequencies, or by selection favouring one combination of alleles over another (Falconer, 1996).

Here, linkage disequilibrium is used (for the first time, as far as the author is aware) to measure niching. Although there are many standard techniques for measuring genetic diversity, often a unimodal (normal) distribution of alleles is assumed, whilst allele associations are ignored. Linkage disequilibrium is an exception: particularly sensitive to allele associations and ignoring underlying genetic variation, linkage disequilibrium is an excellent measure of niching. Assuming  $n$  loci and  $N$  individuals, linkage disequilibrium is calculated as follows (taken from Barton & Gale, 1993). For each individual  $i$ , let  $z_i$  be the proportion of loci at which allele  $A$  (allele "1" for the coin toss game) is found. For each loci  $l$ , let  $p_l$  be the proportion of individuals at which allele  $A$  is found. Then, let mean and variation of  $z_i$  be  $\bar{z}$  and  $var(z)$ , respectively. Also, let mean and variation of  $p_l$  be  $\bar{p}$  and  $var(p)$ , respectively. Then, *average pairwise linkage disequilibrium*,  $\bar{D}$ , is:

$$var(z) = \frac{1}{2n} \left( \bar{z}(1 - \bar{z}) - var(p) \right) + \frac{1}{2} \left( 1 - \frac{1}{n} \right) \bar{D} \quad (6.2)$$

### 6.2.3 Results

Figure 6.6 displays four typical graphs resulting from the four test conditions of the coin toss game. The large graphs display the number of H-alleles each individual contains each generation; with hosts shaded dark and parasites light. The smaller graphs, below, show the linkage disequilibrium of genotypes within each population; roughly approximating the amount of niching. Clearly, both resource sharing and reduced parasite virulence affect the dynamics of coevolution.

Under condition one—Maximum virulence with no resource sharing, i.e., typical coevolutionary optimisation (figure 6.6, top left)—the system exhibits cycling. After the initial generations, hosts may begin to recruit more H-alleles in order to defeat parasites tossing Heads. However, as parasites counter-adapt, by recruiting more T-alleles, they increase the likelihood of tossing Tails. In response, hosts appear with a greater proportion of T-alleles, with the entire population eventually switching strategy, in order to concentrate on winning the T-allele half of the game. Subsequently, parasites again regain the *upper-hand* by tossing Heads more frequently, and so on. Under these conditions, the coin toss game is inherently easier for the parasite population. Hosts find it difficult to be

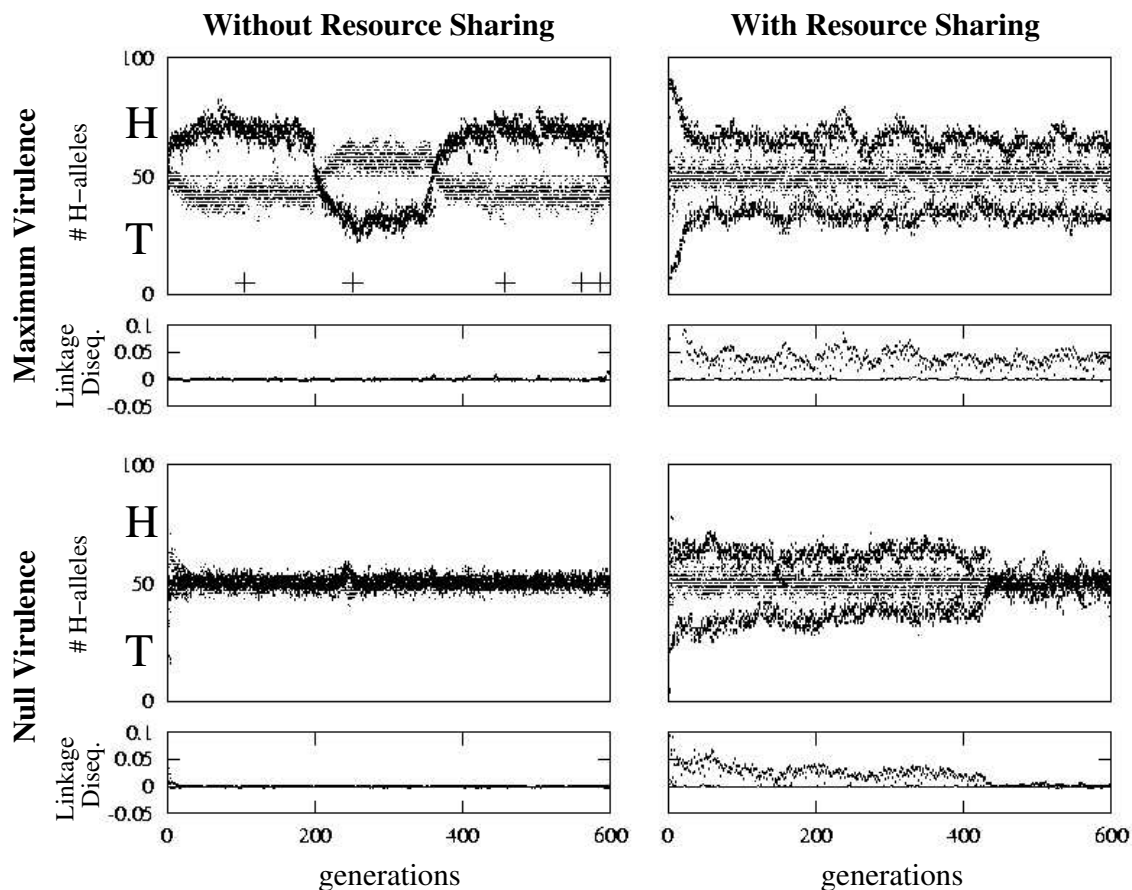


Figure 6.6: Typical coevolution of the coin toss game. Maximum virulence without resource sharing produces cycling as hosts repeatedly alternate strategy (top left). Maximum virulence with resource sharing leads to mediocre stability with half the host population focusing upon each strategy (top right). Null virulence without resource sharing encourages hosts to become generalists, capable of matching some parasites whatever result is tossed (bottom left). Null virulence with resource sharing initially pushes hosts into two specialist niches, before funnelling the population into generalists (bottom right).

successful generalists—incapable of matching parasites along *both* dimensions—and are encouraged to become brittle *specialists*. As a result, maximally virulent parasites win the majority of competitions and occasionally win *all* competitions, resulting in disengagement (indicated by crosses).

Under condition two—Maximum virulence with resource sharing (figure 6.6, top right)—the system reaches mediocre stability, or collusive disengagement. At the beginning of the run, hosts immediately niche into two groups, each specialising on one half of the matching game. In order to be as unpredictable as possible, parasites tend towards 50% of each allele—any deviation from this distribution will be punished by one of the specialist host niches. At this mediocre equilibrium the host population *as a*



*whole* achieves roughly 50% victories over parasites; but each individual host is extremely vulnerable to parasites tossing the opposite result. In contrast, parasites tend to become maximally unpredictable, tossing Heads and Tails with roughly equal probability.

Under condition three—Null virulence ( $\lambda = 0.5$ ) without resource sharing (figure 6.6, bottom left)—the system stabilises with generalist hosts. After the initial generations, the host population settles into generalist strategies capable of matching some parasites whichever result is tossed. Moderate virulence ensures that parasites are rewarded when occasionally matched, thus allowing hosts to succeed without having to concentrate on winning one half of the coin toss game. It should be noticed that reducing virulence *does not* result in host-parasite collusion, which would tend to result in homogeneous parasites; the simplest to match. Rather, parasites remain challenging and unpredictable. Any deviation from a 50% mix is quickly punished by the generalist hosts. As such, both hosts and parasites engage in competition in the most *difficult* regions of space. This is equivalent to discovering the “play random” strategy in scissors-paper-stone, or a generalist immune system capable of defeating a wide range of intruders.

Under condition four—Null virulence ( $\lambda = 0.5$ ) with resource sharing (figure 6.6, bottom right)—the system initially achieves mediocre stability (collusive disengagement), before hosts become generalists, strongly engaged with parasites. Early in the run resource sharing encourages the host population into two niches, each concentrating on one half of the game. In this way, the system reaches mediocre stability with hosts and parasites sharing victories. However, unlike condition two, mediocre stability does not persist. Recall that Null virulence encourages parasites to achieve a win-rate half that of the highest scoring parasite. This scheme lures parasites away from the mediocre equilibrium at which they achieve a 50% win-rate. As parasites become more easily matched, they reduce the pressure upon hosts to concentrate on one half of the matching game. In this way hosts are steered towards more generalist strategies. Hosts engage parasites in a difficult region of space, unattainable without a reduction in parasite virulence. This “funnelling effect”—hosts initially diversify into specific niches before funnelling into generalists—suggests great scaffolding potential for more complex problems, introducing otherwise unattainable intermediate strategies. As far as the author is aware, this is the first time such an effect has been observed.

Results clearly demonstrate that imposing reduced virulence on parasites alters coevolutionary dynamics in a fundamentally different way to that achieved by resource sharing. Whilst resource sharing encourages *within-population* genetic (and phenotypic) diversity, observable as niching in the host population, reduced virulence encourages diversity in *relative fitness* (i.e., a between-population phenomenon).

See section 7.5 for further (CIAO plot) visualization, and discussion, of these results.

## 6.2.4 Caring versus Sharing

Resource sharing encourages a population to diversify into separate niches, thus reducing the likelihood of *over-focusing*. In this way, coevolutionary cycling may also be avoided. However, niching may produce mediocre stability (or collusive disengagement) whereby niches *share* success. In contrast, reducing virulence does not encourage intra-population diversity and, rather, encourages engagement: the extent to which coevolving populations *interact*.

Resource sharing adds a second layer of coupling between conspecifics. In addition to the standard competition that conspecifics experience—striving to beat more opponents than each other—they are forced to *share* their success with one another. This encourages individuals to beat different opponents; i.e., to be different from one another. Niching results from this additional intra-population coupling.

In contrast, reducing parasite virulence increases *inter-population* coupling: it ensures that individuals in one population *care* about the success of individuals in the other. In particular, through attempting to achieve moderate success, parasites care about the variation in relative fitness achieved by their opponents: they are selected to cause a range of scores in their opponents. However, this is not achieved through niching, or genetic diversity *per se*. Rather, it is a direct consequence of the moderation that maintains engagement.

It is true that increased genetic diversity has *some* relationship with coevolutionary engagement. If genetic diversity reduces to zero, populations will disengage (individuals will achieve equivalent scores).<sup>3</sup> However, the converse is not true. Genetic diversity *does not* ensure engagement. Both populations may feature a diverse array of phenotypes, yet still suffer disengagement if each and every phenotype in one population defeats each and every phenotype in the other. Indeed, periods of disengagement often increase genetic diversity through random drift, without necessarily increasing engagement. While this coevolutionary coupling (engagedness) is affected by genetic diversity (and noise, sampling error, etc.), it is not determined by it. Further, niching may directly lead to collusive disengagement: with each niche sharing results, individuals once again become indiscriminable.

These considerations ensure that reducing virulence and resource sharing are *complementary*, rather than *exclusive*, tools. It is not necessary to choose one over another.

---

<sup>3</sup>Attentive readers of a biological disposition may note that some genetic diversity, but no phenotypic diversity, will also lead to disengagement. However, the argument presented here does not benefit from a discussion of genotype-phenotype mappings.

Indeed, the greatest success may result from using diversity *and* engagement maintenance techniques in conjunction (as observed by Juillé & Pollack, 1998a, 1998b).

## 6.3 Sorting Networks

Investigations in the simple numbers and coin toss domains have gathered evidence to support the hypothesis that reducing parasite virulence can encourage coevolutionary engagement. However, it is not obvious that these results generalise to other, more complex, domains. In order to address this issue, reduced parasite virulence is applied to the coevolution of minimal length sorting networks; a domain that has historically attracted interest from the coevolutionary computing community (refer to section 3.3). The aim is to design the shortest fixed *network of comparisons* that can sort, into numerical order, any input list containing a specific number of elements. Comparisons exist in the form *if  $a > b$  then swap, else do nothing*.

### 6.3.1 Background

Conveniently, fixed length sorting networks may be represented graphically. Figure 6.7, left, displays a 5-comparison network that can sort all 4-element input lists. List elements move through the network from left to right. At each comparison (represented by a vertical bar) if the top integer is larger than the bottom, then list elements are swapped, else they are left alone. Here, the sorting network compares elements (1,2), (0,3), (0,1), (2,3) and (1,2) in 4 parallel steps: the network has *depth* 4. Whilst each step must be performed in strict order, multiple comparisons within one step—e.g., (0,1) and (2,3) at step 3—can be performed in either order, or at the same time. When hard coded as a parallel algorithm, the depth of sorting networks is very important (though not considered in this section).

Figure 6.7, right, shows an example list [4,9,8,2] being input into the network. Since  $9 > 8$ , step 1 swaps elements 1 and 2, leaving [4,8,9,2]. Step 2 swaps integers 4 and 2, leaving [2,8,9,4]. Since  $2 < 8$ , step 3 leaves elements 0 and 1 in position, but swaps integers 9 and 4, leaving [2,8,4,9]. Finally, step 4 swaps integers 8 and 4, outputting a fully sorted list [2,4,8,9]. Following this procedure, *any* 4-element list will be sorted.

In this section, 45-comparison networks are evolved to sort 13-element input lists.




Figure 6.7: Left: standard graphical representation of a 4-element sorting network with 5 comparisons. Right: example list input [4,9,8,2] is sorted, producing output [2,4,8,9].

### 6.3.2 Set Up

For historical reasons<sup>4</sup>, 13-element sorting networks have been chosen to coevolve. Rather than attempt an assault on the minimal length record, however, the aim is to observe the difference in performance that varying parasite virulence induces.

The model used here is loosely based upon Hillis' original scheme, however, as Hillis was primarily interested in optimisation, several changes are implemented. Hillis utilised very large populations (of the order of  $10^6$  individuals), sexual recombination, and seeded initial populations with the *butterfly*<sup>5</sup> (Rosin, 1997). As the primary aim is not to find an optimal network, Hillis' (1990) model is simplified by using asexual populations of reduced size, and by initialising individuals at random (i.e., no butterfly seeding). In order to test the effects of spatial embedding, runs are performed under two conditions: with spatial embedding and without spatial embedding. Under both conditions, hosts and parasites evolve on a toroidal grid, with exactly one host and one parasite occupying each location and playing each other.

**With Spatial Embedding** For each grid location,  $g$ , a tournament is played between individuals in the Von Neumann neighbourhood (i.e., centre square and 4 nearest neighbours; North, East, South, West). The highest scoring host and parasite leave (perhaps mutated) progeny in the centre square,  $g$ , the following generation.

**Without Spatial Embedding** For each grid location, a host and a parasite are chosen from two tournaments of 5 individuals each, randomly selected from across the grid—highest scoring individual always chosen—each leave (perhaps mutated) progeny at  $g$ .

Host networks consist of 45 pairs of integers, with each pair representing list elements to be compared and, if necessary, swapped. Host mutation occurs at each loci with probability 0.02, producing a random integer in the range 1 to 13. Parasites each contain 40 unsorted 13-element lists. In order to preserve lists as permutations of the integers 1 to

<sup>4</sup>The shortest network currently known for this problem was discovered using a coevolutionary algorithm (Juillé, 1995).

<sup>5</sup>The first 32 exchanges (the butterfly) in Green's 60-comparison network are known to sort the vast majority of inputs. As such, Hillis seeded every initial individual with the butterfly in order to encourage optimisation.

13, parasite mutation consists of swapping two elements of the list—this occurs at each loci with probability 0.02. Both parasite and host population sizes are identical, with each host attempting to sort exactly one parasite—that which shares the same grid location. Without spatial embedding, therefore, parasites and hosts are essentially paired at random each generation, however, with spatial embedding, specific host and parasite lineages are more likely to meet repeatedly over multiple generations. Both hosts and parasites are asexual.

A host is rewarded with fitness proportional to the number of parasite lists that are completely sorted. Reciprocally, parasites are rewarded for possessing lists that remain unsorted. Maximum, Moderate and Null virulence and the Phantom parasite were each tested. In conjunction, Random parasites with no heredity were also tested so as to compare the coevolutionary results of each virulence scheme against what is effectively *standard* evolution. In order to collect accurate statistics, 30 runs were carried out for each condition with population sizes ranging from 25 to 225. An advantage of this problem domain is that an absolute, objective fitness measure of hosts is possible: networks are given every possible input to sort<sup>6</sup> with the percentage of correctly sorted lists determining an absolute fitness performance. A host network that can sort 100% of all possible inputs is an optimal network.

Under each condition, performance is compared by calculating the mean of the absolute performance of the best individual discovered each run. It should be noted that (like the numbers game, section 6.1) the coevolutionary system has no knowledge of this *absolute* fitness performance; it is merely a way to record comparable results.

### 6.3.3 Results

#### 6.3.3.1 Non-Spatial Model

Figure 6.8 displays the results of coevolving sorting networks *without* spatial embedding. The five conditions are labelled Maximum, Moderate, Null, Phantom and Random. Each graph shows the mean (over 30 runs) of the current best-so-far network performance.

Without spatial embedding, the Moderate and Null conditions significantly outperform Maximum, Phantom and Random conditions, particularly when population size is small. However, this performance difference reduces as population size is increased, suggesting that—since disengagement becomes progressively less likely as population sizes increase—reducing virulence improves performance by diminishing the effects of disen-

---

<sup>6</sup>The set of input lists can be calculated efficiently by using the zero-one principle: ‘a network can sort every input list if and only if it can sort every binary input list’ (Knuth, 1973).

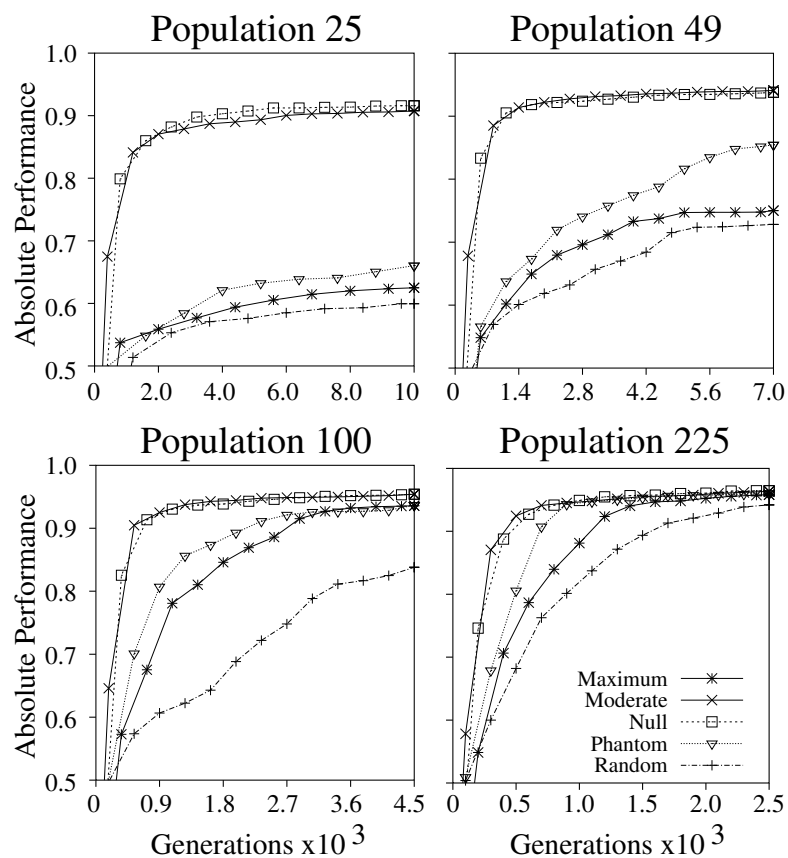


Figure 6.8: Coevolution of minimal length sorting networks with global (non-spatial) interactions. As population size increases, the number of generations graphed is reduced to depict approximately equivalent lengths of units of computational time.

gement (rather than for some other reason).

Maximum does not perform significantly better than Random until population size reaches 100. Considering the weak selective pressure induced by Random parasites this is a poor result, again suggesting that Maximum virulence produces long periods of disengagement. Indeed, upon scrutinising individual runs (not shown), this can be observed.

### 6.3.3.2 Spatial Model

When spatial embedding is implemented (figure 6.9) Moderate and Null conditions significantly outperform Maximum, Phantom and Random conditions across all population sizes. Once again, this can be understood in terms of disengagement. Spatial embedding exacerbates the asymmetrical advantage favouring parasites—as host and parasite lineages repeatedly meet over many generations, parasites are able to specialise against specific host weaknesses (overfitting or overspecialising). As spatial embedding allows only geographically local interactions, this effect is largely independent of population

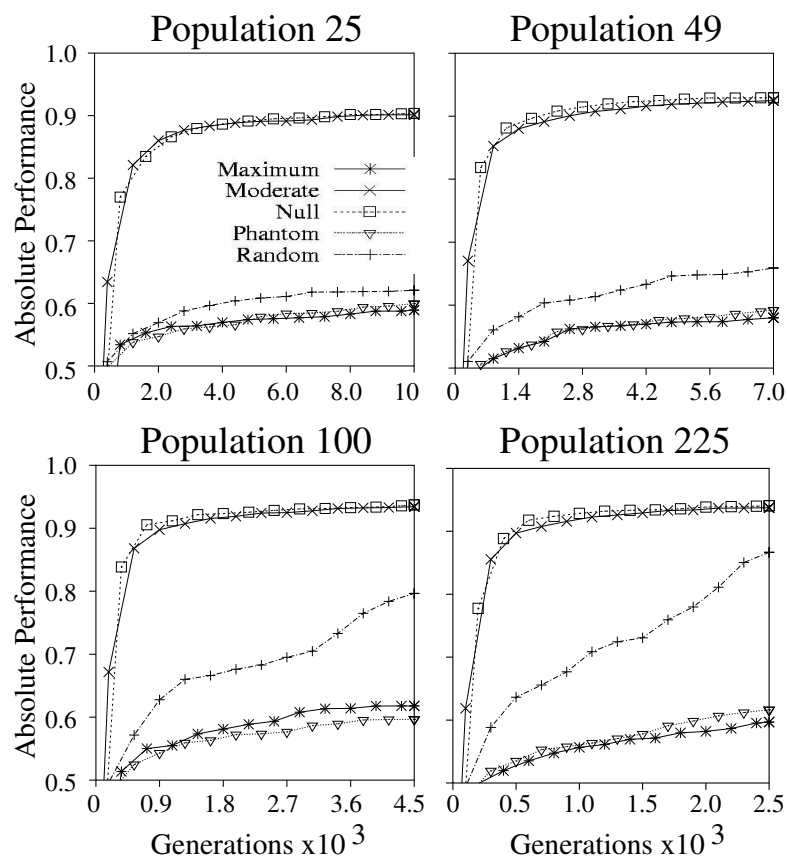


Figure 6.9: Coevolution of minimal length sorting networks with local (spatial) interactions. As population size increases, the number of generations graphed is reduced to depict approximately equivalent lengths of units of computational time.

size. Whilst the Maximum condition ensures that it is very difficult for hosts to engage, Moderate and Null encourage engagement by stopping parasites from overfitting host idiosyncrasies, thus resulting in improved performance.

These results are suggestive, but are certainly not conclusive. Evolving minimum-comparison sorting networks is a complex problem domain, making it difficult to analyse. However, we would expect the list-sorting problem to exhibit the kind of asymmetry that led to disengagement in the counting-ones problem. Being a challenging set of lists is much easier, at least at the outset of coevolution, than being an accomplished list sorter. Could this asymmetry account for the relatively slow progress made by conventional coevolution? If this type of initial asymmetry is a feature of many coevolutionary optimisation problems, reducing parasite virulence in some way could turn out to be an approach with wide application.

## 6.4 Dynamic Parasite Virulence

### 6.4.1 Rationale

Until this point, reduced parasite virulence has been implemented with fixed parameter value  $\lambda$ . However, in order to drive a system to higher levels of performance, it is likely that varying  $\lambda$  over the course of a run will be necessary (see discussion, section 6.1.3.2). Virulence must simultaneously remain low enough to maintain engagement, but be high enough to encourage progress by imposing strong selective pressure. As the system moves through genotype space, the best “compromise” value for  $\lambda$  is likely to vary.

In domains more “realistic” than the numbers game and the coin toss game, problem asymmetry will not tend to be constant over the course of evolution. Initially, one population may enjoy an advantage over the other in terms of the ease with which successful mutant counter-adaptations can be generated. Subsequently, this asymmetry may wax and wane, or even reverse: a population of near-optimal sorting algorithms may enjoy this type of advantage over their parasitic competitors. Under these conditions, in order to maintain an ideal balance between engagement and selective pressure,  $\lambda$  values may need to be constantly varied (for each population). The obvious way to address these concerns is to produce an algorithm that *dynamically adapts* virulence levels in response to the current state of the coevolutionary system.<sup>7</sup>

Here, a methodology for producing a process capable of dynamically moderating the virulence function of parasites is developed. Still in its formative stages, this preliminary work is primarily a proof of concept.

### 6.4.2 Measuring Engagedness

To create a dynamic virulence function, a quantifiable measure of disengagement is necessary: it is not sufficient to qualify systems as *slightly* or *fully* engaged. This section introduces a method for defining the disengagement of a system—a *metric* of disengagement—to be used to vary parasite virulence dynamically. Ideally, any metric should be informative, robust and computationally inexpensive.

Coevolutionary disengagement occurs when there is a lack of discrimination between the relative fitness scores of individuals. Thus, it is easy to intuit that discrimination is a good *indicator* of engagedness. Since discrimination is easy to calculate, it is used here to vary virulence.

---

<sup>7</sup>Initial work on a manually-guided virulence algorithm is underway (refer to section 7.6); allowing human controllers to vary the value(s) of  $\lambda$ , as required, during the course of a run.



Discrimination may be simply measured by the standard deviation of (relative) fitness scores within a population: more variation implies more discrimination, and thus more engagement. However, standard deviation is sensitive to the distribution of fitness scores—multimodal distributions are not accurately reflected—and does not offer a “maximum” value that can easily equate to full engagement. Thus, for the purpose of varying virulence, standard deviation is *not* an ideal statistic. For this reason, a novel (and slightly contrived) method of approximating engagement is developed here.

Since the relative fitness of an individual is measured as the success against a sample of opponents, there are a discrete number of possible fitness levels that each individual,  $g$ , can achieve. In the numbers game, for example, each individual can win, lose or draw against each opponent in the sample,  $\mathcal{S}$ . Thus, there are  $1+2|\mathcal{S}|$  possible discrete values of fitness. By considering each of these values as an equivalence class—individuals in each class have equivalent fitness—diversity within a population can be estimated, simply, by calculating the proportion of classes containing at least one member.<sup>8</sup> At least one equivalence class must always be non-empty. In a disengaged population every individual has the same relative fitness—all individuals are equivalent—and engagement is 0. If individuals are found in all equivalence classes, or sparsely distributed amongst many classes, then the population is maximally discriminated. Let  $C$  be the total number of equivalence classes,  $c$  be the number of non-empty classes, and  $N$  be population size, then engagement  $\varepsilon$  is calculated as:

$$\varepsilon = \begin{cases} \frac{c-1}{C-1} & \text{if } C < N \\ \frac{c-1}{N-1} & \text{otherwise} \end{cases}$$

For example, if 8 out of 11 possible classes are non-empty (population size 20), then  $\varepsilon = 0.7$ ; if 6 out of 40 are non-empty (population size 21) then  $\varepsilon = 0.25$ .

### 6.4.3 Implementation

Below, two dynamic virulence functions are introduced: Dynamic 1 and Dynamic 2. Whilst both utilise  $\varepsilon$  (described above) to vary virulence, Dynamic 1 is designed to tend towards lower levels of virulence,  $\lambda$ , than Dynamic 2.

#### Parasite Fitness:

$$f(x, \lambda) = \frac{2x}{\lambda} - \frac{x^2}{\lambda^2}$$

---

<sup>8</sup>If there are more equivalence classes than there are individuals, then diversity is calculated as the proportion of non-empty classes to population size.

**Dynamic Virulence 1:**

$$\lambda = \begin{cases} 0.5 & \text{if } \varepsilon \leq 0.5 \\ \varepsilon & \text{otherwise} \end{cases}$$

**Dynamic Virulence 2:**

$$\lambda = \begin{cases} 0.5 & \text{if } \varepsilon < 0.25 \\ 2\varepsilon & \text{if } 0.25 \leq \varepsilon \leq 0.5 \\ 1.0 & \text{otherwise} \end{cases}$$

The virulence functions Dynamic 1 and Dynamic 2 each produce a set of curves which range from Null to Maximum (refer to figure 5.2). However, Dynamic 2 will often result in higher virulence,  $\lambda$ , than Dynamic 1. Consider Dynamic 1: if  $\varepsilon \leq 0.5$ , a Null virulence curve is produced; if  $\varepsilon \geq 0.5$  then a  $\lambda = \varepsilon$  virulence curve is produced. Consider Dynamic 2: if  $\varepsilon \leq 0.25$ , a Null virulence curve is produced; if  $0.25 \leq \varepsilon \leq 0.5$  then a  $\lambda = 2\varepsilon$  curve, between Null and Maximum, is produced; if  $\varepsilon \geq 0.5$  then a Maximum virulence curve is produced.<sup>9</sup>

It is assumed that dynamic virulence will aid coevolutionary progress by reducing virulence to counteract disengagement, when necessary, whilst increasing virulence to encourage periods of rapid evolution at other times. In the following section these assumptions are tested.

**6.4.4 Investigations**

To investigate dynamic parasite virulence the numbers game and sorting network domains are utilised once again. To aid comparison with previous results (sections 6.1–6.3), experiments are repeated under the same conditions.

**6.4.4.1 The Numbers Game**

Following the set up of section 6.1.2, figure 6.10 displays three typical runs of the numbers game with parasite mutation bias  $B_{par} = 0.75$  and mutation rate  $m = 0.03$ . The same initial conditions (random seed) as figure 6.1 are used for each run. When employing maximally virulent parasites (figure 6.10, left) the system twice disengages. However, swap-

<sup>9</sup>Although the curves chosen for Dynamic Virulence 1 and 2 are piece-wise linear, there is no reason to expect this to be optimal. Indeed, it is quite likely that it is not. Virulence may well be more sensitive to some values of  $\varepsilon$  than others, making non-linear curves more practical. However, since this work is only intended as a conceptual proof, the simple (piece-wise) curves described above are sufficient.




Figure 6.10: The Numbers Game dynamic virulence.

ping to Dynamic 2 at generation 250 circumvents the second period of disengagement (figure 6.10, centre). Utilising Dynamic 2 from the start (figure 6.10, right) maintains system engagement throughout. In comparison to Moderate virulence (see figure 6.1), Dynamic 2 maintains a higher level of performance, occasionally reaching optimality. Dynamic 1 (not shown) maintains a slightly lower level of performance, more similar to that achieved by Moderate virulence (figure 6.1). Hence, the dynamic virulence functions are both capable of reducing disengagement in the numbers game.

To test whether dynamic virulence reduces disengagement more generally, the experiments of section 6.1.3.1 were repeated. Figure 6.11 displays the result of 30 runs using host and parasite tournament size  $T_h = T_p = 5$ , mutation rate  $m = [0.005, 0.050]$  and parasite bias range  $B_{par} = [0.5, 0.99]$ . Across all levels of asymmetry, Dynamic 1 results in very little disengagement. Comparing figure 6.3, Dynamic 1 is outperformed *only* by Null virulence when asymmetry is very large (parasite bias and mutation rate is high); when asymmetry is very low, Dynamic 1 outperforms Null virulence. More likely to encourage higher levels of virulence,  $\lambda$ , Dynamic 2 does not perform as well as Dynamic 1 when asymmetry is large. Comparing figure 6.3, Dynamic 2 performs slightly better than the Phantom Parasite, but worse than Moderate virulence. Both dynamic functions perform much better than Maximum virulence.

In general, the dynamic functions appear able to reduce disengagement in the numbers game. Can they also encourage near-optimal solutions with high values of  $n$ ?

In section 6.1.3.2, the effects of virulence upon the performance of coevolution was tested by ensuring populations maintained engagement. This was achieved by making the host population an *exact replica* of the parasite population each generation. Results suggested that Maximum virulence, reaching a mean level of 99.1% (30 runs), produces the highest performance. Performance was shown to reduce with virulence.

To test the dynamic virulence functions, the experiment in section 6.1.3.2 was re-




Figure 6.11: The average of 50 runs (each 600 generations) of the numbers game using dynamic virulence.

peated. Impressively, Dynamic 1 averaged  $99.2 \pm 0.33\%$  (95% confidence limit) and Dynamic 2 averaged  $99.0 \pm 0.76\%$ ; much higher than that achieved by Moderate ( $94.8 \pm 1.50\%$ ) and Null ( $75.1 \pm 2.59\%$ ) virulence. This is easily explained. In each run, populations are manipulated to remain highly engaged (large  $\varepsilon$ ). Thus, for prolonged periods, the dynamic functions exhibit Maximum virulence. The resulting performance is very high.

In the numbers game, at least, the dynamic virulence hypotheses appear correct. Dynamic 1 and 2 produce efficient coevolution and near optimal performance by utilising high virulence during times of high engagement, and low virulence, when required, to stop disengagement.

#### 6.4.4.2 Sorting Networks

The previous section demonstrated that the dynamic functions exhibit the properties required of a self-moderating virulence function. To test the generality of these observations, this section investigates Dynamic 1 and 2 in the more complex domain of sorting networks. In section 6.3, it was demonstrated that reducing the virulence of parasites can aid in the coevolutionary optimisation of sorting networks. To test Dynamic 1 and 2, the experiments of section 6.3 were repeated under the same conditions. The average results of 30 runs are shown in figure 6.12.

Under both global (non-spatial; figure 6.12, top) and local (spatial; figure 6.12, bottom) interactions, the dynamic virulence functions achieve high performance across all population sizes. Comparing these results with figures 6.8 and 6.9, it can be observed that Dynamic 1 and 2 perform at a level statistically indistinguishable from Null and Moderate virulence (but much higher than Maximum, Random and the Phantom Parasite). Under the assumption that higher virulence encourages a greater level of performance—if en-

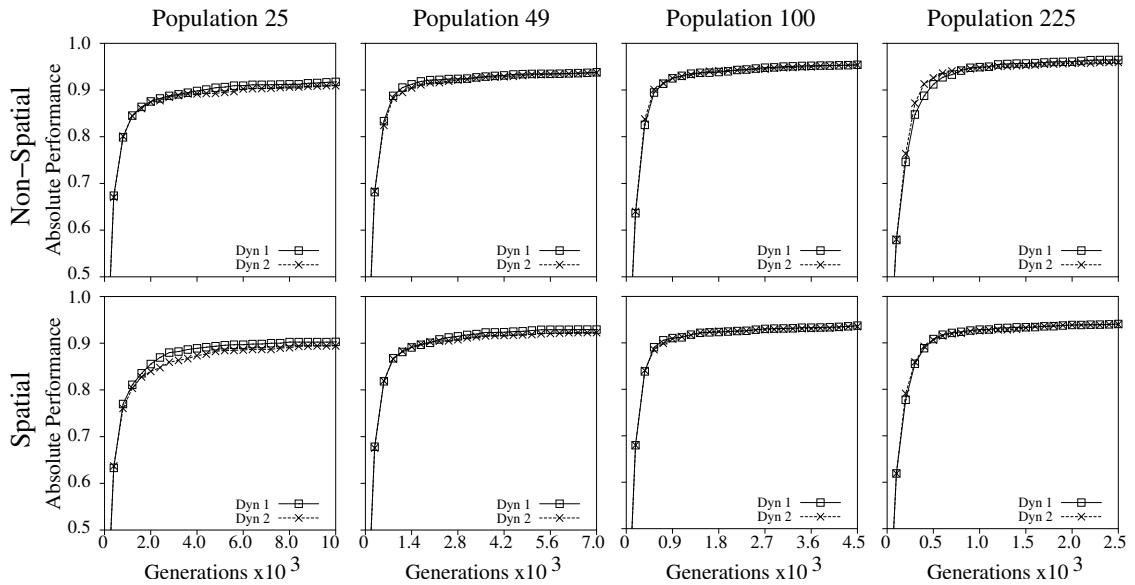


Figure 6.12: Results of coevolving sorting networks. The dynamic virulence functions, Dynamic 1 and Dynamic 2, perform equivalently to Moderate and Null virulence (compare figures 6.8 and 6.9), but much better than Maximum, Random and the Phantom Parasite.

gement is maintained—this is unexpected.

Throughout all runs, for the majority of generations,  $\varepsilon$  varies between 0.2 and 0.5. At such times, Dynamic 1 exhibits Null virulence; thus, it is little wonder that Dynamic 1 and Null virulence perform equivalently. However, the behaviour of Dynamic 2 is less easy to understand. Regularly exhibiting high or Maximum virulence, it is surprising that Dynamic 2 does not outperform Dynamic 1: if anything, it performs slightly worse (not statistically significant). The reason for this is not clear; perhaps periods of high virulence push the system to near-disengagement, thus removing useful host diversity.

Given the preliminary nature of these experiments, the sorting network results are not entirely disappointing; much work is left to be done in this area. Firstly, the current process of measuring engagement,  $\varepsilon$ , is fairly crude; it is likely that much better measurements exist. Secondly, translating  $\varepsilon$  to  $\lambda$  via the dynamic virulence functions can also be improved. Rather than calculate  $\lambda$  as a linear (truncated) function of  $\varepsilon$ , perhaps dynamic virulence should be sensitive to particular engagement values (a non-linear function), or take the engagement of previous generations into consideration (evolutionary momentum). Finally, it is possible that the sorting network domain is not suitable for dynamic virulence; other domains of study should be considered (see section 6.5).

In conclusion, this section has succeeded in demonstrating that virulence can be dynamically manipulated during run time on the basis of system engagement. As proof of

concept, this is a valuable preliminary result.

## 6.5 Summary

The reduced parasite virulence technique (introduced in chapter 5) has been shown to combat disengagement in a number of domains, under a spectrum of conditions. When compared to the Phantom Parasite (introduced in section 3.4.5)—the only domain independent competitor—reduced virulence performed favourably in all investigations. Further, reduced virulence has been shown to complement standard genetic diversity maintenance techniques (such as resource sharing; introduced in section 3.4.2); it is *not* an alternative tool. Finally, a first attempt at dynamic virulence has been introduced, demonstrating that virulence can be automatically varied with system engagement.

The problem domains explored in this chapter were primarily chosen for simplicity and resulting ease of analysis. However, the No Free Lunch theorem (Wolpert & Macready, 1995) warns that reduced virulence will be effective on a particular subset of problems only. Which problems are good candidates for the approach?

Firstly, the reduced virulence technique is likely to be of use in problems where a reasonable (but not necessarily optimal) solution is required under strong constraints of limited time or computational resources (e.g., dynamic load allocation across a telecommunications network) since reducing parasite virulence can accelerate rapid initial progress towards high-quality solutions despite small population sizes.

Secondly, where the character of the problem is constantly changing and demands constant evolutionary change in the solution population (such as maintaining a strong immune response in a changing environment) reducing virulence is likely to improve performance by encouraging and maintaining engagement.

Finally, for problem spaces where small genetic changes often give rise to qualitative changes at the phenotype level and consequent discontinuous jumps in fitness (e.g., chess strategy) the reduced virulence technique may improve performance by promoting re-engagement, when it occurs, rather than actively resisting it in the manner of traditional coevolution.

Coevolutionary algorithms are very valuable and versatile tools, yet there remain an ensemble of problems restricting their successful application. Hopefully, this chapter has demonstrated that the introduction of reduced virulence has moved coevolutionary computation (in general) one step closer to realising its potential.

# Chapter 7

## Visualizing Coevolutionary Dynamics

---

Evolutionary computation (EC) concerns the study of high-dimensional, time-varying systems that exhibit complex dynamics on a number of time scales and levels of organisation. Yet, before the potential of evolutionary algorithms can be fully realised, it is imperative that their behaviour is well understood. Informative and intuitive visualizations offer a method of generating insight into evolutionary dynamics, but there has currently been relatively little work in this area. While many idiosyncratic graphing techniques have been developed for particular systems, there remains a reliance on simplistic plots of summary statistics: visualizations that inherently disguise much system complexity.

Perhaps the only visualization technique designed specifically for the study of coevolutionary dynamics is the CIAO plot (Cliff & Miller, 1995, refer to section 3.4.1). As a result, CIAO plots regularly appear in the coevolutionary literature. In this chapter, an investigation of CIAO plots reveals the existence of previously unrecognised ambiguities in interpretation. To remedy this, it is recommended that CIAO plots be used as part of a complementary suite of visualization tools, rather than in isolation. This multi-visualization method is used to analyse a biological simulation of *E. coli*.

Following the investigations of CIAO plots, a first attempt towards coevolutionary steering—an *interactive* visualization tool—is developed, demonstrating the possibility of manipulating system parameters to guide/steer coevolving populations during run time. The prospective future of (co)evolutionary steering for education, research and application, is very exciting.

## 7.1 Background

### 7.1.1 EC Visualization

A diverse array of EC visualization techniques exist: some represent the evolving populations, others the character of the problem being solved; some present information at the end of a run, others show the development of a system during run time; some present loci- or gene-level representations, some individual-level variables, and others population-level statistics (e.g., Kapsalis & Smith, 1992; Dabs & Schoof, 1995; Collins, 1998; Bosman, 1999; Wu, De Jong, Burke, Ramsey, & Grefenstette, 1999; Pohlheim, 2001; Hart & Ross, 2001). However, of this wide variety, few have been published outside of technical reports.

The range of visualization tools employed in evolutionary computing includes standard techniques such as various types of multi-dimensional scaling (e.g., Spears, 1994; Collins, 1999; Pohlheim, 1999), Sammon mapping (Sammon, 1969; Dzwiniel, 1994; employed by Dybowski, Collins & Weller, 1996) and quadcodes (Li & Loewn, 1987; independently developed by Collins, 1997; Shine & Eick, 1997; Wiles & Tonkes, 2002), and entirely novel approaches developed specifically for dealing with EC issues such as techniques for representing genotypic changes over evolutionary time (Wu et al., 1999; Hart & Ross, 2001). However, none of these platforms or techniques has achieved significant penetration in the EC community, as of yet.

One stream of visualization research utilises EC theory by suggesting what type of data will be informative, in what way, and in which situations. In addition to generating useful visualization tools, this approach is intended to progress EC theory by providing a richer insight into the behaviour of evolutionary systems. However, studies that have combined EC visualization and theory are still relatively rare: Cliff and Miller (1995, see below) proposed a visualization technique for detecting cycling in coevolutionary systems; Harvey and Thompson (1996) explored the use of various visualization methods in order to explain the role of neutral ridges in the evolutionary search space; Bedau and Brown (1998) visualized an evolutionary activity metric; and Bullock (2001) used visualization to demonstrate the biases inherent in a range of mutation operators.

With the considerable potential to build on these initial studies, the future of EC visualization, in general, seems bright. In the following section, the particular challenges associated with visualizing coevolutionary systems are discussed.



## 7.1.2 Visualizing Coevolutionary Systems

In an effort to either engineer artificial adaptive systems or better understand natural systems, researchers have employed *coevolutionary* algorithms as design tools and simulation models. However, while this research remains promising, it has raised several problematic issues. Perhaps the most pressing of these concerns our ability to understand the dynamics of coevolutionary systems (refer to section 3.4).

In particular, there are inherent difficulties associated with detecting coevolutionary progress. Often, it is virtually impossible for an observer to know whether populations are improving over time. A short-term improvement in fitness, relative to contemporary competitors, does not necessarily lead to long-term improvement in some objective sense, since it is possible for coevolutionary systems to cycle: coevolving populations may follow a repeating sequence of adaptive transitions.

In standard evolutionary algorithms, because the fitness of an individual is measured against a static function, continual progress can be detected as an improvement in fitness over time. However, this is not true of coevolutionary systems, where an individual's fitness is measured *relative* to its contemporary opponents. When this type of relative fitness measure is plotted against time, individuals from different generations are being compared using a metric which is itself varying unpredictably over time, since the opponents against which they were assessed will not, in general, have been the same. These considerations suggest that such plots are not merely “difficult to interpret”, but are effectively meaningless. Unfortunately, this ensures that detecting the occurrence of cycling in either natural or artificial systems is problematic, since an external observer cannot determine, on the basis of relative fitness measures, whether a coevolutionary system is progressing, cycling, or drifting randomly.

In an attempt to circumvent this problem, Cliff and Miller (1995) proposed the CIAO plot (refer to section 3.4.1) as a visualization tool for detecting coevolutionary progress. At the conclusion of a coevolutionary run, CIAO plots can be constructed by pitting the elite (i.e., best-scoring) individual from every generation against the elite opponent from each ancestral generation and plotting the results as shaded cells in a matrix. In this way individuals are directly assessed against the ancestors of their opponents. If, over many generations, most individuals can beat their ancestral opponents, the matrix will exhibit a consistent gradation in shading from dark cells at the origin to light cells at the leading edge. This pattern suggests that there has been continual progress over the course of coevolution, since individuals from later generations are outperforming their ancestors. By contrast, coevolutionary cycling manifests itself as a diagonal “banded” pattern (see section 7.2).

In coevolutionary research, CIAO plots are a widely accepted problem-independent visualization technique with few alternatives. However, the studies reported here explore whether they are as easy to interpret as has previously been implied. In the following sections, it is demonstrated that CIAO plots can be misleading even in a simple coevolutionary domain. It is shown that coevolutionary cycling can fail to produce a characteristic banded CIAO plot, resulting instead in a “tartan” pattern: a class of CIAO plot that is commonly reported but has received little attention. In section 7.3, it is established that tartan-like CIAO plots can result from cyclic coevolution that is *irregular*, or from random drift. It is thus argued that CIAO plots are vulnerable to ambiguity and that, as a result, their use should be accompanied by more problem-specific analysis.

## 7.2 CIAO Plots

Appearing in Lewis Carroll’s *Through The Looking Glass*, the Red Queen must continually run in order to maintain her position. No matter how fast she moves, the surrounding landscape always keeps up with her.

van Valen (1973) made an analogy between the Red Queen and biological coevolution after discovering a surprising trend concerning the probability of species extinction. After analysing huge data sets collected across a wide range of biological taxa, van Valen (1973) noticed that, counter to intuition, “all groups for which data exist go extinct at a rate that is constant for a given group”. Assuming that species evolve in a relatively static environment, one would expect beneficial adaptations to accumulate over evolutionary time, enabling progressive generations to be better equipped at defending against extinction. To explain his findings, van Valen proposed a new evolutionary law of extinction, with the Red Queen Hypothesis as its central tenet: “biotic forces provide the basis for self-driving . . . perpetual motion of the effective environment”. Any beneficial adaptation by a particular species is inevitably detrimental to other species inhabiting the same effective environment. Coevolutionary forces will, in turn, select for specific counter-adaptations in these species. In this way, adaptive advantage is continually eroded. In Red Queen fashion, species continually struggle to maintain their relative fitness.

One might expect Red Queen dynamics to drive a run-away process of continual counter-adaptation—a coevolutionary *arms race*. Such arms races are considered a profound force driving evolutionary adaptation in the natural world, and have sometimes been characterised as a source of strong selection for novel adaptations capable of accelerating evolutionary progress (Dawkins & Krebs, 1979).

In order to take advantage of arms-race dynamics, the field of evolutionary compu-

tation has seized upon coevolution as an attractive alternative to standard evolutionary optimisation. Living up to expectation, artificial coevolution has had success in several domains (refer to section 3.3 for a review). However, as an optimisation technique, coevolutionary search suffers from the relative nature of fitness assessment—coevolutionary systems can be difficult to drive in a consistent, objectively “progressive” direction.

There are several ways in which an arms race can unfold. One side may “win” the race, discovering an adaptation to which there is no available counter-adaptation. If coevolution is within a species, evolutionary stasis may be reached, but during between-species coevolution the disadvantaged species may be driven extinct. Alternatively, one population may temporarily win by “outstripping” the other to such an extent that the opponents are no longer discriminated by selection: whilst a counter-adaptation may exist, it has yet to be discovered (such “disengagement” is discussed further in section 3.4.5 and throughout chapter 6). Finally, an arms race may cycle (see section 3.4.3), as populations follow repeated trajectories through strategy space, discovering strategies that enjoy only a temporary advantage over their opponents. From the perspective of coevolutionary optimisation, cyclic coevolution has gained the most attention. Since these cyclic trajectories waste computational resources, it would be useful to detect (and ultimately prevent) their occurrence; or, at least, to realize that the optimisation problem as-stated has a possibly unanticipated set of unstable equilibria with no escape trajectory.

In general, one would like to know how coevolution is progressing. Has there been steady and continuous improvement, or transient bursts of progress amidst long periods of stasis? Is the system cycling, and, if it is, are the cycles regular? Frustratingly, the Red Queen renders standard approaches to answering these questions—plotting individual fitness as it changes over time—obsolete (refer to section 3.4.1).

Consider a predator-prey arms race. As predators *improve* over evolutionary time, one might expect them to catch more prey. The prey, however, are improving too. “There is no general reason to expect the average success of animals at out-running or out-witting *contemporary* enemies, victims, prey or competitors, to improve over evolutionary time” (Dawkins & Krebs, 1979, italics added). This is a direct consequence of the Red Queen. An improvement in any one species is countered by each coevolving species, resulting in a deterioration of the effective environment (van Valen, 1973).

However, one might expect a progressive coevolutionary arms race to result in an advantage for modern predators and prey over their *ancestral* adversaries. In nature, it is difficult to perceive how a competition across evolutionary time may arise, without the use of cloning or time-travel. Yet, in simulation it is quite feasible to carry out such ancestral opponent competitions. Cliff and Miller developed this technique during a series of papers




Figure 7.1: Schematic of a “Current Individual versus Ancestral Opponents” (CIAO) plot. At each square  $(x, y)$  the result of competition between the A-elite from generation  $y$  and the B-elite from generation  $x$  is plotted as a grey-scale value, with increasingly heavy shading representing an increasing margin of victory in favour of A. The leading diagonal (far-right) plots A-elite( $g$ ) against B-elite( $g$ ) for all generations,  $0 \leq g < N$ . The diagonal immediately to the left plots each generation’s A-elite against the *previous* generation’s B-elite. Horizontal rows plot the results of A-elite( $y$ ) against *all* ancestral B-elites( $0 \dots y$ ) (adapted from Cliff & Miller, 1995).

in which they attempted to coevolve pursuer and evader strategies in continuous-time neural-network controllers (Miller & Cliff, 1994a, 1994b; Cliff & Miller, 1995, 1996).

### 7.2.1 Plotting CIAO Contests

Once a coevolutionary run has terminated, ancestral opponent contests (competitions against ancestral opponents) are carried out between the highest scoring individual of each generation: the “elite”. The elite of population A (the A-elite) is pitted against the elite of population B (the B-elite) in a series of contests. For each generation  $g$  of coevolution, the A-elite( $g$ ) competes with the B-elite of the current, and each ancestral, generation. Hence, A-elite(5) plays B-elite(5), B-elite(4),  $\dots$ , B-elite(0). The resulting scores of each contest are normalised and converted into grey-scale values and plotted on a 2-dimensional grid: the greater the victory in favour of the A-elite, the heavier the shading of the relevant matrix cell (see figure 7.1, adapted from Cliff & Miller, 1995).

### 7.2.2 Progress and Cycling

Along with the invention of CIAO plots, Cliff and Miller left a lasting legacy: the idealised plot. Idealised CIAO plots demonstrate the patterns that they predicted one would find, given either perfect continual coevolutionary progress, or perfect coevolutionary cycling. Example schematics are shown in figure 7.2. Continuous progress appears as continuous

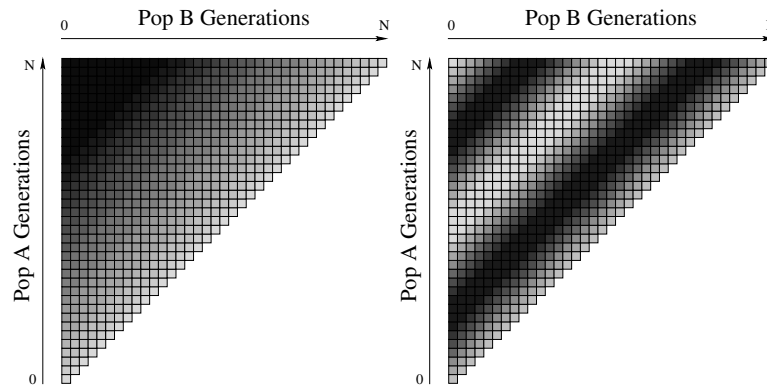


Figure 7.2: Idealised CIAO plots showing smooth progress (left) and cycling with regular period (right). With each horizontal row depicting the performance of one A-elite against each ancestral B-elite, smooth progress produces a gradation in intensity from dark to light: A-elites beat ancestral opponents, the more ancestral the greater the victory (i.e., the further left, the darker the cell). In contrast, regular cycling manifests as diagonal banding. Whilst the latest A-elites beat recent ancestral B-elites (diagonals to the right are dark), they perform less-well against more ancestral B-elites (middle diagonals are light), but outperform even earlier ancestors (diagonals further left are dark), etc.

diagonal *gradation* from dark to light across the plot. Each individual beats its ancestral opponents; the earlier the ancestor, the greater the victory. In contrast, coevolutionary cycling produces diagonal *banding*. Whilst recent ancestral opponents are easily beaten, elite opponents from a few generations earlier have the upper hand. As the contests span increasing periods of evolutionary time, competitive advantage repeatedly transfers between novel and ancestral populations: the system is cycling.

### 7.2.3 Unpicking Tartan CIAO Plots

Since Cliff and Miller (1995) first introduced CIAO plots for pursuit and evasion, they have become widely accepted as a standard tool for visualizing coevolutionary progress. As such, CIAO plots have received extensive use in the fields of evolutionary robotics, (Floreano & Nolfi, 1997a, 1997b; Nolfi & Floreano, 1998; Stanley & Miikkulainen, 2002), the games of Go (Lubberts & Miikkulainen, 2001), 3-D Tic-Tac-Toe and Nim, (Rosin & Belew, 1997), and in the coevolution of string generators and predictors (Ficici & Pollack, 1998b).

However, the results obtained in these relatively complex domains rarely resemble the idealised plots of Cliff and Miller. Even discounting noise, real CIAO plot visualizations are often qualitatively different from the ideal schematics. Many CIAO plots exhibit a tartan pattern—a patchwork of unpredictable lines and rectangles across the plot




Figure 7.3: Schematic of a “tartan” CIAO plot. Periods of competitive stasis (blocks of uniform shade) are separated by sharp transitions in competitive advantage (block boundaries). Such an irregular pattern is referred to as “tartan” throughout this paper.

(figure 7.3 presents a hypothetical example). Of 22 plots found in the literature, 10 are tartan in nature, 8 show progress (smooth gradation) and 4 show no progress (a largely homogeneous plot). There are no examples of regular banding. In a tartan plot, blocks of uniform shade represent periods of stasis bounded by adaptive innovations. Consider Rock–Paper–Scissors, for example (see below): if the elites in both populations play Rock from generation 50 onwards, until population A’s elite adopts Paper at generation 55, and the B-elite adopts Scissors at generation 60, then the resulting CIAO plot will display a grey rectangle with sides  $5 \times 10$  generations, cornered (50, 50).

Tartan CIAO plots have sometimes been interpreted as supporting evidence for cycling (Floreano & Nolfi, 1997a, 1997b; Nolfi & Floreano, 1998). The principle reason appears to be that, unlike the idealised progressive CIAO plot, there is little correlation between grey-scale value and time. It is difficult to imagine how a jumble of dramatic jumps between winning and losing can occur across evolutionary time in a continuously progressing coevolutionary system. How such a pattern relates to cycling, however, appears unclear. It shall be shown (below) that although tartan plots of this kind are consistent with coevolutionary change that is cyclic but irregular, it is dangerous to infer irregular cycling from them uncritically. The following section demonstrates that tartan plots are, in fact, ambiguous with respect to showing irregular cycling and/or random drift.

## 7.3 Rock–Paper–Scissors: An Intransitive Domain

Here, coevolution is explored in a simple domain specifically chosen to exhibit cyclic coevolutionary trajectories.<sup>1</sup> To better understand the underlying coevolutionary dynamics, several different visualizations are presented. However, whilst each individually offers insight into the underlying coevolutionary dynamics, one cannot interpret them with confidence until information from the whole suite is considered: this is particularly true if some are ambiguous.

This section considers the coevolution of two populations of strategies for the simple parlour game Rock–Paper–Scissors (RPS). In each bout of the game, two players simultaneously choose one of the three possible moves (Rock, Paper, Scissors), the winner being decided according to an intransitive superiority relationship (Rock blunts Scissors, Scissors cuts Paper, Paper covers Rock; draws are shared). Although the global optimum for either player is to play each move with equal probability, any deviation from this strategy on the part of either population encourages a complementary counter-deviation.

If one population is biased in favour of playing Paper, for example, the other will benefit from playing Scissors more often than one third of the time. As the frequency with which Scissors is played increases, the first population is under selection pressure to increase the frequency with which its members play Rock. In response, the opponent population will tend to favour Paper. In this way, while both populations are continually favouring offspring better able to compete against their current opponents, in the longer term, the populations are continually and repetitively cycling through a sequence of globally sub-optimal strategies as they seek to exploit the temporary biases of their opponents.

### 7.3.1 Study 1: A Simple Encoding

In this baseline study the inherent intransitivity of the Rock–Paper–Scissors (RPS) game, in combination with the smooth fitness landscape resulting from a simple genetic encoding, produces regular coevolutionary cycling. The visualizations that reflect this include CIAO plots that closely resemble the idealised banded plot (figure 7.2). For the most part this section introduces fairly straightforward results and visualizations which are intended to contrast with the more complex results from the subsequent studies.

For the simple encoding, each genome consists of three positive integers that sum to 100. Each integer represents the probability of choosing one of the three game moves, Rock (R), Paper (P), or Scissors (S). For example, genome  $\{50, 0, 50\}$  represents an indi-

---

<sup>1</sup>Rock–Paper–Scissors is the simplest game to exhibit underlying intransitivity— $a > b > c > a$ —likely to induce cycling.

vidual that chooses to play Rock and Scissors with equal likelihood, but will never play Paper. Genome  $\{33, 33, 34\}$  represents a near-optimal individual that plays each move randomly with almost equal probability.

During reproduction, mutation occurs with probability 0.1 per locus. A mutated gene is incremented by an integer drawn at random from the uniform distribution  $[-30 \dots 30]$ . Following mutation, the genome is *normalised* to once again sum to 100.

Two reproductively isolated (initially randomised) populations, each containing 20 RPS players, are coevolved for 256 generations. Each generation, every individual is pitted against every member of the opponent population. Each game consists of 10 bouts, with each individual choosing R, P, or S probabilistically on the basis of their genes. Throughout a generation, an individual accumulates a score from its 200 competitions. Individuals reproduce asexually, with tournaments biasing selection in favour of high-scoring individuals. The winner of each randomly assembled 5-member tournament is chosen to reproduce.

### 7.3.1.1 Feature Detection

To assist the detection of features within the CIAO plots reported in this chapter, a three stage processing of the raw images is employed. The method chosen (described below) is generalizable and easy to implement. Further, Marr (1982) has suggested that it may have biological implications as a model for the very first stages of visual processing. For further details, one should refer to Marr (1982, chapter 2).

1. **Gaussian Blur:** A 2-dimensional Gaussian filter (radius  $r$ ) is initially applied to the image; effectively removing structures smaller than the standard deviation of the Gaussian distribution. As it is smooth and omnidirectional, a Gaussian distribution is appropriate for blurring as it is unlikely to introduce structure that was not present in the original image.
2. **Laplacian of Gaussian (LoG):** The Laplacian ( $\nabla^2$ ) is an isotropic second-order differential operator that can be used to detect intensity changes in a blurred image, as seen at the scale of the particular Gaussian employed (determined by the standard deviation  $\sigma$ ). The Laplacian operator is used to locate the zero-crossings of an image—changes from dark to light, or light to dark—thus generating a contour map.

In mathematical notation, the blur of an image function  $I(x, y)$  with a Gaussian function  $G$  is denoted by  $G * I$ , read as “ $G$  convolved with  $I$ ”. The Laplacian of



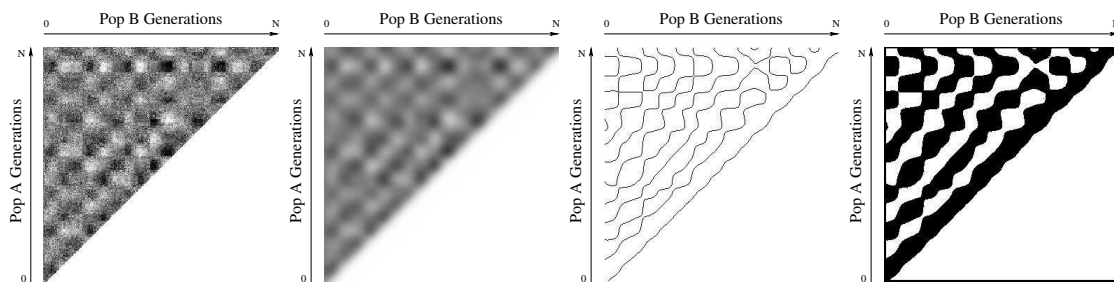


Figure 7.4: Study 1 results. CIAO plots depicting one representative run over  $N = 256$  generations. From left to right: (1) Diagonal banding in the raw data plot suggests regular cycling; (2) A large Gaussian ( $r = 20$ ) removes fine detail; (3) LoG ( $\sigma = 8$ ) produces a contour map; (4) The image is binarized. The fully processed plot displays clear diagonal banding.




Figure 7.5: Study 1 results. Event plot for a single representative run, labelling the generations in which the A-elite plays Paper or the B-elite plays Rock with probability 0.66 or greater. The repeated “Paper follows Rock” event sequence suggests regular cycling.

this is denoted  $\nabla^2(G * I) = (\nabla^2 G) * I$ , read as “the Laplacian of Gaussian (LoG) convolved with  $I$ ”.

3. **Binarize:** Finally, the direction of each zero-crossing is highlighted by “binarizing” the image. Areas with negative values are coloured black whilst areas with positive values are left white.

Whilst stages 1 and 2 of the feature detection process naturally combine into one operation, throughout this chapter they will be performed independently. This enables the result of blurring to be viewed before zero-crossings are detected. Throughout,  $r$  and  $\sigma$  values are chosen through trial and error: values are reduced to detect smaller features, for larger features they are increased.

### 7.3.1.2 Results

Throughout this section, all data is the result of one representative coevolutionary run. Figure 7.4 displays four CIAO plots. From left to right, plot 1 displays the raw data set, with subsequent CIAO plots presenting the same data after each stage of the image

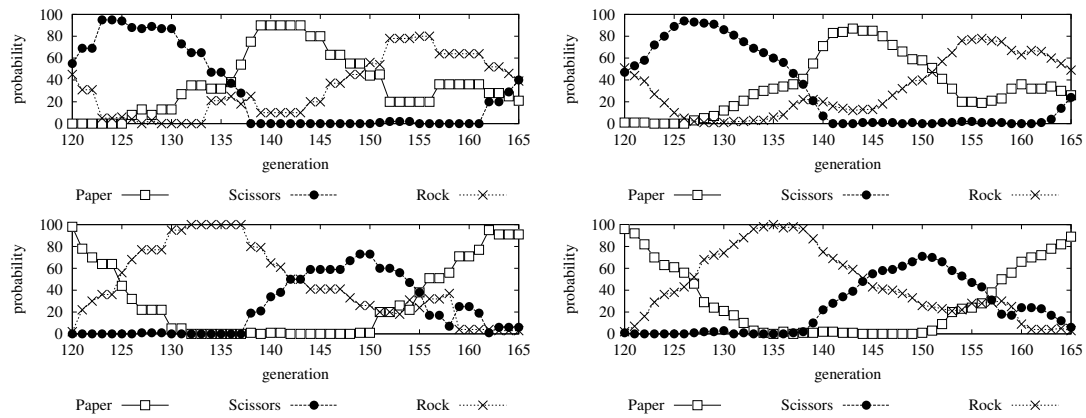


Figure 7.6: Study 1 results. Graphs plotting the probability values of playing each RPS move. Top-left: the A-elites. Top-right: the mean probabilities for all individuals in population A. Bottom-left: the B-elites. Bottom-right: the mean probabilities for all individuals in population B. The system completes one full cycle between generations 120 and 165.

processing routine described in section 7.3.1.1. In plot 1, regular patterns and possible diagonal banding are discernable, potentially indicating regular coevolutionary cycling. After blurring the image with a large Gaussian of radius 20 pixels, small features are removed and diagonal banding becomes more pronounced (plot 2). A contour map is produced using a Laplacian of Gaussian (LoG) with standard deviation 8 pixels: diagonal banding is clear (plot 3). Plot 4 shows the result of binarizing the contour map in order to highlight the direction of any zero-crossings. The regularity and clarity of diagonal banding in the fully processed plot is striking. A comparison with the idealised banded CIAO plot of figure 7.2 allows plot 4 to be easily classified as demonstrating coevolutionary cycling with regular period: the feature detection routine has clearly enhanced the interpretability of the original CIAO plot data.

However, to have confidence in the nature of cycling, further visualizations were constructed. For this purpose, an “event plot” is used to highlight each occurrence of specific events during coevolution: evidence for (or against) cycling is generated by the resulting sequence of events. Assuming events are chosen adequately, regular cycling should manifest as a repeated sequence of events with fixed period. The event plot of figure 7.5 records the occurrence of events associated with the play of the best scoring individual of each population: the A-elite and B-elite. The plot is marked for each generation in which the A-elite plays  $P$ , or the B-elite plays  $R$ , with probability  $p > 0.66$ . Figure 7.5 clearly shows a repeated sequence. As expected, once B-elites regularly play Rock, A-elites quickly evolve to play Paper. After approximately 40 generations, the system completes a full cycle with B-elites once again predominantly playing Rock. The coevolutionary sys-

		To		
		R	P	S
From	R	0	9	1
	P	1	0	9
	S	8	1	1

Figure 7.7: Study 1 results. State transition table. Repeatedly following state transitions “from Scissors to Rock”, “from Rock to Paper” and “from Paper to Scissors”, the system is clearly cycling in a regular manner.

tem appears to be following a cyclic trajectory with regular period. In general, if mutation rate *or* selection pressure (in the form of tournament size) increases, the cycle period decreases.

In figure 7.6, the elite (i.e., best-scoring) and average (the mean of all individuals) strategy of each population are depicted for a representative period of evolutionary time. The vertical axis represents the probability of playing each of the three possible moves. One full cycle takes place during the 45 generations depicted. At generation 120, the elite member of population B (bottom-left) always plays Paper (genotype  $\{0, 100, 0\}$ ). Within 5 generations, Scissors dominates the elite strategy in population A (top-left). As a counter-adaptation, population B converges on Rock, which in turn drives the elite of population A to play Paper with 90% probability (genotype  $\{10, 90, 0\}$ ) by generation 140. This completes a half-cycle. Scissors becomes dominant for population B elites around generation 150, followed by Rock for population A, and finally the cycle is completed around generation 165 as the elite strategy of population B returns to Paper (genotype  $\{3, 91, 6\}$ ). The mean population strategies (right) lag their elite counterparts slightly, but demonstrate the same trend. Notice that each population is able to counter-adapt to its opponents in a smooth, regular manner.

Figure 7.7 tabulates the state transitions of study 1; classified as follows. Each population enters a new state,  $m$ , the *first time* an elite of that population plays move  $m \in \{R,P,S\}$  with probability  $p > 0.66$ . Figure 7.7 records the transitions of each population *in response to the other*. Thus, the first time an elite of *either* population plays move  $m \in \{R,P,S\}$  with  $p > 0.66$ , the system enters initial state,  $m$ . A transition occurs the next time the *opposing* population enters a new state. For example, if the A-elite of generation 1 plays Rock with  $p > 0.66$ , then population A (and the system as a whole) enters state P. Suppose the B-elite of generation 5 plays Paper with  $p > 0.66$ , then population




Figure 7.8: Complex genotypes encode a deterministic finite state machine (FSM) that governs player behaviour. State transitions are dictated by an opponent’s play (represented by transition labels). Here we see that machine *A* beats machine *B* over a five-bout contest. The final score is 3:2 in favour of *A*. As both machines have returned to their initial states, further bouts will produce repeated play.

*B* enters state *P* and a transition “from *R* to *P*” is recorded. The next transition occurs when population *A* enters a new state: if the *A*-elite of population 7 plays Scissors with  $p > 0.66$ , then there is a transition “from *P* to *S*”. Figure 7.7 displays the state transitions of an entire run. As expected, the system repeatedly follows the state transitions “from Scissors to Rock”, “from Rock to Paper” and “from Paper to Scissors”, with little deviation. Clearly, the system is cycling in a regular and predictable manner.

In summary, study 1 has (predictably) demonstrated that simple regular cycling can manifest itself on CIAO plots as cyclic banding, closely resembling the idealised banded CIAO plot in figure 7.2. In conjunction, it has also been demonstrated that greater insight into cyclic coevolutionary behaviour can be obtained through the additional use of alternative, perhaps problem-specific, visualizations. In the following section, the complexity of the genetic encoding is increased to demonstrate that, whilst cycling may persist, the ease with which it is visualized is considerably reduced.

### 7.3.2 Study 2: A Complex Encoding

In this study a more complex genetic encoding scheme is employed. Although the strategic structure of the RPS game remains unchanged, by increasing the complexity of the strategies available to players, more complex courses of coevolutionary adaptation are available. It is intended that this new encoding will influence the system’s coevolutionary dynamics, ensuring that the search for counter-adaptations more closely resembles that experienced in more realistic models of coevolutionary competition.

Here, genomes code for a deterministic finite state machine (FSM). Example genomes are presented in figure 7.8. The start node defines the choice of play during the first bout of a contest. In figure 7.8, both individuals initially play Scissors. From each node there are

exactly three out-edges (which may be self-connecting), with each transition representing an opponent's move. In figure 7.8, since both individuals begin by playing Scissors, each follows the state transition  $S$ .  $A$  returns to the start node (0) and continues to play Scissors whilst  $B$  transfers to a new node (1) with state Paper. In the second bout of the game,  $A$ 's Scissors beats  $B$ 's Paper.  $A$  follows the  $P$  transition associated with  $B$ 's play and once again returns to state  $S$  (node 0). In response to  $A$ 's Scissors,  $B$  follows the  $S$  transition to a new node (2) with state Rock.  $B$  is victorious in bout three, as Rock beats Scissors.  $A$  moves to state  $P$  (node 1) and  $B$  remains at state  $R$  (node 2). Figure 7.8 details the full results of a five-bout contest.

Each node contains a unique integer identifier, a play state, at least one in-edge and three out-edges, each associated with an opponent's play and each connecting to a legal node. Both populations are initialised with random, self-connected, single-node FSMs. During reproduction, three mutation operators are employed (described below).

**Node Mutation:** With probability 0.03 per genome, a node is either added or removed.

Any edges previously connected to a removed node become self-connections. FSMs were constrained to have between one and one hundred nodes. In practice, however, FSMs rarely grew above 10 nodes.

**State Mutation:** With probability 0.02 per node, a node state is mutated to one of the other two states, chosen at random.

**Edge Mutation:** With probability 0.02 per edge, edges are mutated by randomly changing the node to which the edge is an in-edge.

Note that the deterministic FSM encoding employed in this study (and study 3) does not allow individuals to reach the global optimum (available in study 1) of playing each move randomly with equal probability, irrespective of opponent play: to achieve this a non-deterministic FSM encoding is necessary. However, this is not problematic since a direct comparison between the results of studies 1, 2, and 3 is not an objective of this section. Rather, the impact of regular and irregular cycling on CIAO plot data is of primary concern. It shall be shown that the deterministic encoding used here is a good choice in this respect.

Two reproductively isolated (initially randomised) populations, each containing 25 RPS players, are coevolved for 256 generations. Each generation, every individual is pitted against every member of the opponent population. Each game consists of 10 bouts, with each individual choosing R, P, or S deterministically on the basis of their genes. Throughout a generation, an individual accumulates a score from its 250 competitions.

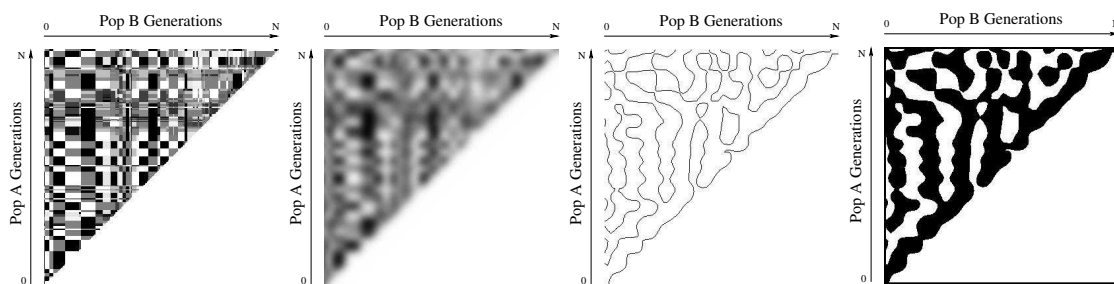


Figure 7.9: Study 2 results. CIAO plots depicting one representative run over  $N = 256$  generations. From left to right: (1) The raw data plot exhibits a tartan pattern that is difficult to interpret; (2) A large Gaussian ( $r = 20$ ) removes fine detail; (3) LoG ( $\sigma = 8$ ) produces a contour map; (4) The image is binarized. Whilst the fully processed plot does not exhibit diagonal banding, some vertical banding is clear. However, one cannot predict regular cycling from this image.




Figure 7.10: Study 2 results. Event plot for a single representative run, labelling the generations in which the A-elite plays Paper or the B-elite plays Rock with effective probability 0.66 or greater. The strong correlation between events suggests (perhaps irregular) cycling.

Individuals reproduce asexually, with tournament selection biasing selection in favour of high-scoring individuals. The winner of each randomly-assembled 3-member tournament is chosen to reproduce.

### 7.3.2.1 Results

Throughout this section, all data is the result of one representative coevolutionary run. Figure 7.9 displays four CIAO plots. From left to right, plot 1 displays the raw data set, whilst subsequent CIAO plots display the same data after each stage of the image processing routine described in section 7.3.1.1. Plot 1 (far left) presents a tartan pattern qualitatively similar to the schematic shown in figure 7.3: irregular blocks of uniform shading indicate periods of competitive stasis separated by rapid transitions in competitive advantage. Plot 1 resembles CIAO plots obtained from more complex (practical) problem domains (e.g., Floreano & Nolfi, 1997a, 1997b). The patchwork effect suggests a lack of progress and may imply cycling. However, the image is difficult to interpret with

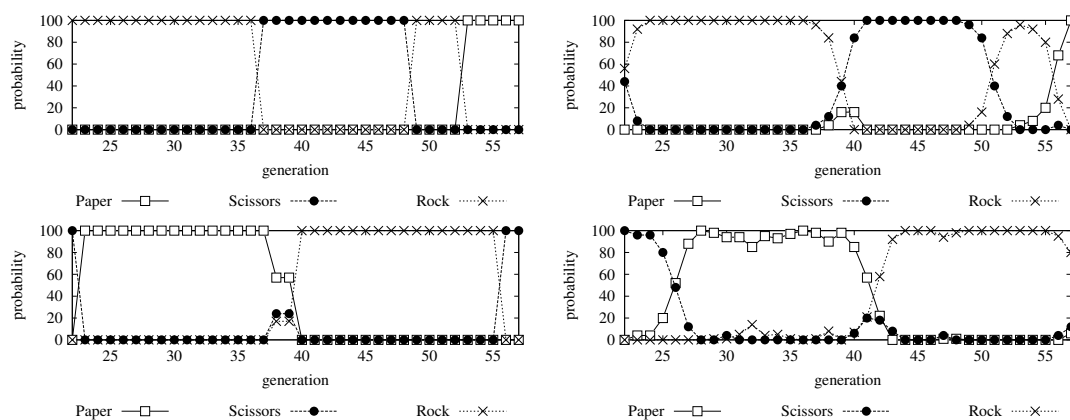


Figure 7.11: Study 2 results. Graphs plotting the effective probabilities associated with playing each move. Top-left: the A-elites. Top-right: the mean probabilities for all individuals in population A. Bottom-left: the B-elites. Bottom-right: the mean probabilities for all individuals in population B. The populations evolve in response to each other, resulting in an irregular cycle.

confidence.

After carrying out the same image-processing process employed in study 1, the plot displays some vertical banding and one diagonal band to the right of the image, but remains difficult to interpret. Whilst the vertical banding may be a result of bias in the original CIAO plot design—adaptive mutations will necessarily appear as banding on the horizontal (if A-elite) or vertical (if B-elite)—the black diagonal band to the right of the plot is potentially interesting. Perhaps the diagonal reflects a competitive advantage enjoyed by each population over its immediate ancestral opponents, or even regular cycling that can be observed over short evolutionary time scales, but that is lost over longer periods. Alternatively, perhaps the diagonal is an artefactual edge-effect of the CIAO plot: a result of the dark plot meeting a white background. In either case, whilst the processed plot has highlighted some potentially interesting structures, ideal banding is not exhibited. Considering the intransitive dominance relationships inherent within the RPS domain, however, one might suspect cycling to be taking place (particularly when the genetic encoding does not allow global stability to be reached). To investigate the underlying coevolutionary dynamics, further analysis is necessary.

The simplicity of the RPS domain enables us to gain insights into its coevolutionary dynamics by direct observation of genotypes and phenotypes as populations change over generations. At the beginning of a run, when each player's FSM has only one node, all strategies are simple. Each player repeatedly chooses the same play irrespective of its opponent's behaviour. Over evolutionary time, more complex strategies arise. These

multi-node FSMs change state in response to an opponent's play. However, as more time passes, simple strategies, such as "always play Scissors", again begin to dominate. This process often repeats several times throughout a coevolutionary run.

For a coevolutionary system to be described as cycling, it must repeatedly enter the same or similar states in the same or similar order. Regular cycling repeats with fixed period, irregular cycling does not. Either class may fail to enter *exactly* the same set of states during each cycle, or maintain *exactly* the same ordering over these states during each cycle. However, each is clearly distinguishable from random drift. Whilst a randomly drifting system may return to previous states over evolutionary time, unless there is a heavy bias influencing the "random" walk (perhaps due to some bias within the genetic encoding or genetic operators, e.g., Bullock, 1999, 2001), the trajectory of change is unpredictable. Any useful method for detecting coevolutionary cycling should be able to distinguish cases of regular or irregular cycling from stochastic repetition that may arise due to random drift. It shall be shown that, when used in isolation, this is something CIAO plots struggle to do.

Choosing suitable visualizations for complex problems is often challenging. The space of variable length FSMs is difficult to represent graphically. In order to re-apply some of the graphing techniques used for the simple encoding scheme, some extra work is necessary.

FSM networks can be re-described simply using three probabilities: the likelihood of playing each of the three possible moves. Since each FSM is a directed graph, an  $m$ -bout contest against an opponent can be described by a path of length  $m$  through the graph. By traversing each of the  $3^m$  possible  $m$ -length paths through the network—the equivalent of playing every possible game—the *effective* probability of a network playing each of the three possible moves can be determined. Since these probabilities are comparable to the three probabilities encoded in the simple genomes of study one, they can be plotted in a similar fashion. However, whilst the effective probabilities of FSMs are comparable to the probabilities of study 1, they are not equivalent: deterministic FSMs do not behave probabilistically but are sensitive to opponent play.

Figure 7.10 shows an event plot for study 2. The plot is marked for each generation in which the A-elite plays P, or the B-elite plays R, with effective probability  $p > 0.66$ . The system repeatedly moves through the same states in a similar sequence. "B-elite plays Rock" is always followed by "A-elite plays Paper". However, on 2 out of 8 occasions, the system enters "A-elite plays Paper" without passing through "B-elite plays Rock". The repeating event plot sequence provides some evidence for irregular cycling: if cycling is occurring, it is not as regular as that exhibited in figure 7.5.



		To		
		R	P	S
From	R	5	12	3
	P	3	1	10
	S	11	2	6

Figure 7.12: Study 2 results. State transition table showing predictable, but irregular, cycling.

Cycling can be directly observed by plotting the effective probabilities of each population's elite and mean strategies during a particular (representative) coevolutionary period (see figure 7.11). Whilst it is not perfectly predictable, sequential cycling can be observed between generations 22-57. The B-elites (bottom-left) and B-means (bottom-right) complete one cycle with the expected sequence SPRS. In response, the A-elites (top-left) and A-means (top-right) also complete one cycle, but with the less predictable sequence RSRP. However, the “unexpected” shift from S to R (rather than directly to P) at generation 49 can be explained by considering the system as a whole at generation 48. With every individual in population A playing S (top-right) and every individual in population B playing R (bottom-right), the system is disengaged in generation 48: every member of population B is beating every member of population A. At this point, a novel mutation of *either* Paper *or* Rock is beneficial to population A: whilst Paper gives victory, Rock allows a draw. Thus, the move from Scissors to Rock is a direct result of disengagement and is not entirely “unexpected” (for a further discussion of disengagement, refer to section 3.4.5). Such “irregular” cycling is confirmed by the state transition table of figure 7.12. Whilst the system is *most likely* to move from R to P, to S, to R, etc., transitions are not as predictable as those of study 1 (figure 7.12).

Comparing the graphs of elite strategies with those of the population mean strategies in figure 7.11 shows that each population exhibits several variable-length periods of strategic stasis. Typically, during these periods, the majority of population A (top) are being beaten by the majority of population B (bottom). Unlike in figure 7.6, population A appears unable to easily discover counter-adaptations to the successful adaptations of population B. This is a result of the deterministic FSM encoding and the limitations of the associated mutation operators. A desired mutation may be difficult to achieve. The initial 57 generations of study 2 are shown in the expanded CIAO plot of figure 7.13. The rapid changes in strategy depicted in figure 7.11 clearly reflect the internal structure of




Figure 7.13: Enlarged CIAO plot depicting the coevolutionary period graphed in figure 7.11. The rapid strategy transitions depicted in figure 7.11 clearly reflect the boundaries in the CIAO plot pattern.

the CIAO plot. The patchwork nature appears because the FSM mutation operators tend to produce either no change in player strategy, or very rapid changes in behaviour.

The different nature of the coevolutionary dynamics generated in study 2 most likely result from the ruggedness and/or neutrality introduced into the fitness landscape by the more complex genetic encoding scheme and associated genetic operators. Single mutations, such as “remove node”, can result in large changes in player behaviour. By contrast, it can often be difficult to make a small change to a player’s behaviour (e.g., increasing the tendency to follow Rock with Paper) without altering several parts of the FSM, which requires several separate mutation events. The simple genetic encoding scheme and associated mutation operator of study 1 ensured that single mutations tended to have modest phenotypic impact, yet most phenotypes were only a few mutations apart. In general, for any population in study 2 some parts of strategy space will be less attainable than others to a larger degree than was the case in study 1. The resulting CIAO plot is difficult to interpret. However, alternative visualizations have shown that cycling is occurring, but that it is irregular.

### 7.3.3 Study 3: Random Drift

Do banded and tartan CIAO plots directly imply regular and irregular cycling, respectively? Can other coevolutionary trajectories lead to banded or tartan CIAO plots? In particular, what kind of CIAO plot is generated by the kind of walk through strategy space produced by random drift? As previously discussed, random drift can generate tra-

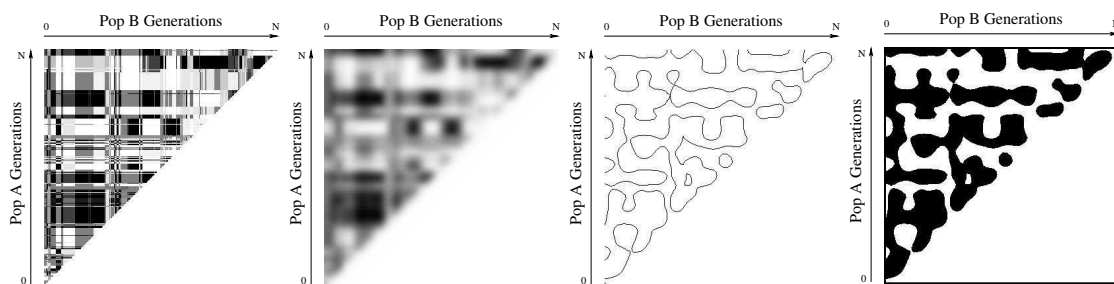


Figure 7.14: Study 3 results. CIAO plots depicting one representative run over  $N = 256$  generations. From left to right: (1) The tartan pattern of the raw data plot is difficult to interpret; (2) A large Gaussian ( $r = 20$ ) removes fine detail; (3) LoG ( $\sigma = 8$ ) produces a contour map; (4) The image is binarized. Whilst there is no diagonal banding, some horizontal banding is clear. However, one can neither predict nor rule out cycling on the basis of the fully processed plot.




Figure 7.15: Study 3 results. Event plot labelling the generations of one representative run in which the A-elite plays Paper or the B-elite plays Rock with effective probability 0.66 or greater. Given the lack of correlation between events, the likelihood of cycling appears small.

jectories that revisit earlier states despite being non-cyclic. Can this kind of repetition be distinguished from regular behaviour using CIAO plots? In order to derive a base-line plot from which CIAOs displaying “interesting” phenomena need to be distinguished, this study attempts to discover the patterns exhibited by CIAO plots portraying random drift.

Random drift is simulated for the same complex encoding scheme and mutation operators employed in study 2 (256 generations with 25 individuals in each population). Evolutionary selection pressures are removed, allowing each individual an equal chance of reproduction irrespective of ability. Individuals reproduce asexually and at random, irrespective of score. Figure 7.14 displays four CIAO plots resulting from typical random drift. From left to right, plot 1 displays the raw data set, whilst subsequent CIAO plots display the same data after each stage of the image processing routine described in section 7.3.1.1. Plot 1 (far left) presents a patchwork pattern of a tartan nature, but less regular than that shown in figure 7.9. The irregular blocks of shading suggest a lack of progress, but may imply cycling of some sort. The plot is difficult to interpret with confidence. After image processing, the binarized plot (far right) displays no diagonal banding.




Figure 7.16: Study 3 results. Graphs plotting the effective probability values of playing each move during random drift. Left: the A-elites. Right: the B-elites. The populations do not appear to evolve in response to each other. Cycling is not observed.

The most prominent features are the horizontal bands across the centre of the image. It is likely that these bands are an artefact of the underlying grid structure of the raw plot, similar to the horizontal banding seen in the binarized plot of study 2 (figure 7.9). The most significant difference between the binarized plots of study 2 and 3 is the diagonal band to the right of figure 7.9 that is missing in figure 7.14. The binarized plot of random drift shows no sign of cycling (diagonal banding) even in the short term. However, whilst the processed plots highlight this feature, they still remain difficult to interpret. Can irregular cycling and random drift be correctly classified on this basis alone?

Figure 7.15 shows an event plot for the random drift of study 3. Once again, the plot is marked for each generation in which the A-elite plays P, or the B-elite plays R, with effective probability  $p > 0.66$ . Unlike the event plot of study 2 (figure 7.10), random drift shows no obvious sequence repetition. Events occur without regularity. The event plot strongly suggests that the system is not cycling.

Figure 7.16 displays the effective probabilities of each population's elite strategies during a particular (representative) period. The state transitions of each population's elites appears random, suggesting that neither cycles. Further, there is no obvious correlation between the two populations: knowing the state of one population does not improve the ability to predict the state of the opponent population. Thus, both populations vary *independently* and at random: it is *not* the case that one population changes randomly, with the other adapting to it in a regular way. This is confirmed by the state transition table of figure 7.17: transitions are highly unpredictable and do not follow the expected pattern “from Rock to Paper to Scissors to Rock”.

Random drift allows populations to return to previously evolved strategies, but not systematically: the system is *not* cycling. Despite this, random drift produces a CIAO plot with a tartan nature, not dissimilar to that produced by the irregular cycling of study 2. While we have not quantitatively measured the differences between the CIAO plots generated in studies 1, 2, and 3, the very fact that we might have to rely upon some

		To		
		R	P	S
From	R	2	11	14
	P	16	7	4
	S	8	10	0

Figure 7.17: Study 3 results. State transition table showing transitions are unpredictable and do not follow the expected pattern of R to P to S to R, etc.

such measure applied to CIAO plots in order to discover what they have to tell us about coevolutionary dynamics is disappointing. To anybody using CIAO plots as their only method of coevolutionary visualization, the results reported here are unfortunate. Not only can coevolutionary cycling result in tartan, rather than banded, plots, but these tartan patterns can also occur in the absence of any cycling at all. If used alone, CIAO plots are potentially ambiguous.

### 7.3.4 Understanding Irregular Coevolutionary Cycling

In section 3.4.3 the prevalence of cycling in natural systems was highlighted with a series of examples demonstrating the ubiquity of the phenomenon. Cycling is *not* confined to “toy” artificial coevolutionary systems, such as those studied here. The three studies carried out above have demonstrated that although one particular kind of coevolutionary cycling is easily detected using CIAO plots, a second class of cyclic behaviour is much harder to detect without resorting to alternative visualizations. Whilst CIAO plots *can* give valuable insights into coevolutionary dynamics, they should preferably be used amongst a suite of techniques in order to enhance their interpretability.

In general, these results contribute to a growing realisation that our understanding of coevolutionary dynamics in artificial systems is far from complete. More specifically, their significance hinges, to some extent, on how important irregular coevolutionary cycling turns out to be. It is argued here that there are good reasons to suppose that this class of dynamic will be more frequently encountered than regular cycling, and that for many kinds of interesting system, when these irregular dynamics are exhibited, they will often be of both theoretical and practical significance.

As evidenced by the different results of studies 1 and 2, as search problems become increasingly complex, their search spaces are increasingly structured by the genetic encod-

ings and genetic operators employed, in addition to the strategic advantages of different phenotypes. This ensures that some (perhaps most) adaptations will only be discovered after a period of evolutionary exploration. The stochastic nature of this exploration coupled with the rugged and/or neutral structure of the search space ensures that the time that this takes is variable. Under such conditions, if a coevolving system finds itself cycling, and hence repeatedly rediscovering the same or similar adaptations in the same or similar order, there is no guarantee that the period of coevolutionary cycling will be constant—indeed it is likely not to be. For these reasons, regular cycling should be regarded as a rarely attainable special case of the more general class of irregular cycling. Perhaps unsurprisingly, of the 22 CIAO plots found in the literature, none display regular cycling (Cliff & Miller, 1995; Floreano & Nolfi, 1997a, 1997b; Rosin & Belew, 1997; Ficici & Pollack, 1998b; Nolfi & Floreano, 1998; Lubberts & Miikkulainen, 2001; Stanley & Miikkulainen, 2002).

Search-space neutrality occurs when many genotypes share the same fitness, perhaps as a result of redundancy in the genetic encoding. A neutral set contains all the genotypes that achieve a particular fitness score, while a neutral network comprises evolutionarily adjacent members of a neutral set. It has been argued that the presence of neutral networks may have profound consequences for the dynamics of evolutionary search. For instance, the neutrality exhibited by natural RNA search spaces has been demonstrated to be of a potentially useful kind, allowing more efficient search (Huynen et al., 1996). More generally, neutrality of the right kind is thought to reduce the chance of premature convergence (Harvey & Thompson, 1996; Barnett, 1998). However, since evolving populations tend to drift at random across neutral networks (but see Bullock, 2002, for analysis of the biases that this drift is subject to) it is difficult to predict how long a population will spend on each one. If a coevolutionary system cycles through a repeated sequence of neutral networks each population will spend some time drifting across each neutral network, before transitioning to the next. As such, although a particular sequence of phenotypes may be generated over evolutionary time, this repetition is unlikely to exhibit a constant period.

Coevolutionary disengagement occurs in a competitive coevolutionary system when one population *outperforms* the other to the extent that different individuals are not discriminated from their contemporaries in fitness terms, i.e., floor or ceiling effects (Watson & Pollack, 2001; Cartlidge & Bullock, 2002, 2004a; also refer to section 3.4.5). For instance, in the coevolution of pursuit and evasion, disengagement could occur if evaders discover a simple hiding strategy that defeats all contemporary opponent pursuers; each hiding evader would score 100%, while all opponent pursuers would score 0%, despite variation in their strategies. When disengagement occurs, selective pressure disappears,

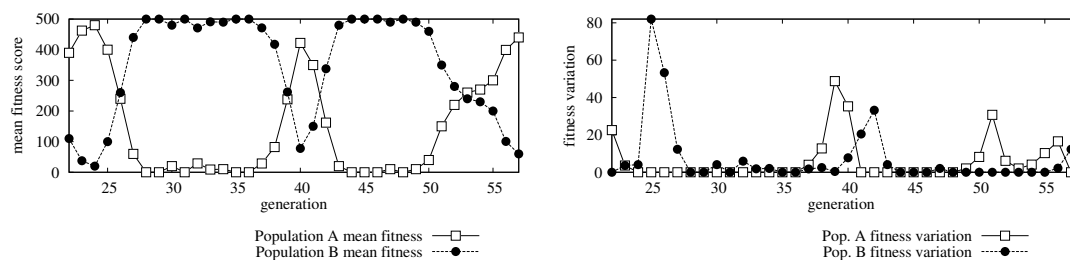


Figure 7.18: Study 2 results. Graph showing the mean fitness (left) and fitness variation (right) within populations A and B. Periods of zero variance signify disengagement. In both populations, there is very little fitness variation during generations 28-36 and 44-49. During these periods, population B is easily outperforming population A. As expected, disengagement coincides with stasis (homogeneous shading) on the enlarged CIAO plot of figure 7.13.

leaving populations free to drift until such time as populations happen to re-engage. Re-engagement takes place when mutant strategies arise that achieve distinctive fitness scores (e.g., a pursuer able to discover hidden evaders, or an evader unable to hide successfully). The time taken for such mutants to arise via neutral drift, and hence the duration of disengagement, is variable for the reasons described above. As a result, a cycling coevolutionary system suffering from disengagement is also likely to cycle with non-fixed period. In short, irregular cycling is likely to be evolutionarily typical because useful evolutionary innovations and counter-innovations are not typically discovered at a constant rate.

The data collected during study 2 exhibits these variable periods of coevolutionary disengagement and neutral drift (in coevolutionary systems the two concepts are closely related). At generation 23, for example, figure 7.11 shows the elite strategy in population one plays Rock, whilst the elite strategy of population two plays Paper. In contrast with the immediate and smooth coevolutionary responses reported for study 1, it is not until generation 37 that the elite strategy of population one discovers the Scissors counter-adaptation. The intervening period is one in which the populations have disengaged, and are drifting across neutral networks of equivalent strategies. Disengagement can be directly observed in figure 7.18, where mean fitness (left) and fitness variance (right) in both populations is plotted over time. During generations 28-26 and 44-49, fitness diversity is very low and occasionally falls to zero (right): here population B is easily outperforming population A (left).

In the following section, CIAO plots are used (for the first time, as far as the author is aware) to analyse simulation data from a biological model: a replication of the Rock–Paper–Scissors *E. coli* experiments of Kerr et al. (2002, refer to section 3.4.3). It is demonstrated that CIAO plots, whilst specifically developed for evolutionary compu-

tation, may benefit evolutionary biology (and the modelling of adaptive behaviour more generally).

## 7.4 Study 4: A Simulation of *E. coli*

Rock–Paper–Scissors cycling has been demonstrated in *E. coli* populations (Kerr et al., 2002). However, practical problems forced experiments on real bacteria to end after short time periods. To collect more “data” (and gain a better understanding of the coevolutionary dynamics) Kerr et al. (2002) ran a simulation. This was shown to behave realistically and suggested that RPS dynamics characterised the system. However, it is argued here that more effective visualization techniques could have afforded greater insight into the underlying dynamics. To test this, the simulation of Kerr et al. (2002) is replicated in this section. The resulting data is then investigated using the visualizations discussed in this chapter (CIAO plots, event plots and probability graphs), demonstrating that the lessons learned in previous studies can be applied to a biological system with potential relevance to evolutionary biologists.

### 7.4.1 Replication

Three strains of *E. coli* exist on a (toroidal) lattice grid of  $100^2$  cells. At the start of each run, each cell is initialised at random (equal probability) to one of four states: occupied by C, S, R, or “empty”. Cells are asynchronously updated. A focal cell,  $c$ , is randomly chosen and updated probabilistically based upon local interactions (the relative states of the 8 nearest neighbouring cells). If  $c$  is empty a bacteria of strain  $i \in \{C, S, R\}$  is chosen to occupy the cell. The probability of choosing  $i$  is given by  $f_i$ , the fraction of the local neighbourhood occupied by each strain. If  $c$  is occupied, the bacteria is killed with probability  $\Delta_i$ . Throughout the simulation,  $\Delta_C$  and  $\Delta_R$  are fixed. However,  $\Delta_S$  varies with  $f_C$  (the fraction of neighbouring C cells) such that  $\Delta_S = \Delta_{S,0} + \tau f_C$ , where  $\Delta_{S,0}$  is the probability of death for S cells with no neighbouring C cells, and  $\tau$  is the toxicity of C cells.

To set up a Rock–Paper–Scissors intransitivity, it is necessary for  $\Delta_{S,0} < \Delta_R < \Delta_C < \Delta_{S,0} + \tau$ , which ensures S displaces R, R displaces C, and C (if sufficiently dense) displaces S. Following Kerr et al. (2002), the following parameter values were chosen:  $\Delta_C = \frac{1}{3}$ ;  $\Delta_{S,0} = \frac{1}{4}$ ;  $\Delta_R = \frac{10}{32}$ ; and  $\tau = \frac{3}{4}$ .

An “epoch” is defined as the mean turnover time across all cells. In the  $100^2$  lattice used here, an epoch occurs every  $10^4$  updates. We consider each epoch to be 1 timestep.






Figure 7.19: Study 4 results. Adopting the visualization scheme employed by Kerr et al. (2002), we plot the locations of each bacterial strain during one representative run of the *E. coli* simulation (toroidal grid size  $100^2$ ). Sensitive (S) bacteria are coloured black, Resistant (R) are dark grey and Col toxic (C) are light grey. Empty cells are white. From left to right: (1) At time step 1000 the three strains coexist across the grid; (2) By time step 1200, C (light grey) have moved into areas occupied by S (black), S have moved to areas originally occupied by R (dark grey), and R have replaced C; (3) S continue to replace R, R replace C and C replace S; (4) By time step 1600, the three strains have returned to roughly the same locations they occupied at time step 1000. One full cycle is complete.

The *E. coli* simulation was run for 5000 timesteps.

## 7.4.2 Results

Adopting the visualization scheme employed by Kerr et al. (2002), figure 7.19 shows four instantaneous “snapshots” taken of the simulation at timesteps (from left to right) 1000, 1200, 1400, and 1600. The state of each cell in the lattice is denoted by its colour. Whilst empty cells are white, those occupied by C, R, or S strains are coloured light grey, dark grey, and black, respectively. It is possible to observe areas of black replaced by light grey, light grey by dark grey, and dark grey by black. The bacteria are following the relationship  $S > R > C > S$  (analogous to RPS). Between timesteps 1000-1600, one full cycle takes place: the system returns to a similar state.

These results qualitatively map those of Kerr et al. (2002) who showed the same succession sequence both in simulation and in real bacteria populations. Whilst the continued coexistence of all three strains and the  $S > R > C > S$  relationship are adequately demonstrated, the visualizations are not wholly satisfying. One cannot determine the nature of cycling from a series of snapshots.

To further visualize the *E. coli* data, the local density of each strain was recorded for one  $5 \times 5$  portion of the lattice. Since the three strains coexist without large fluctuations in global density, plotting CIAO plots for the entire lattice is uninformative. Due to the Rock–Paper–Scissors intransitivity of *E. coli* strains, each local density can be considered

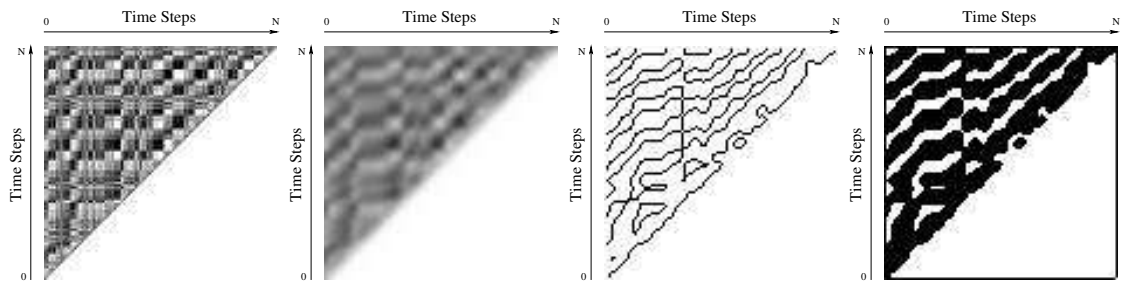


Figure 7.20: Study 4 results. CIAO plots depicting one representative run of the *E. coli* simulation over  $N = 5000$  timesteps (sampled at a resolution of 50 generations). From left to right: (1) Some diagonal banding is observable in the raw data plot; (2) A large Gaussian ( $r = 5$ ) removes fine detail; (3) LoG ( $\sigma = 2$ ) produces a contour map; (4) The image is binarized. The fully processed plot exhibits clear diagonal banding.

equivalent to the probability of playing a particular move in RPS. For example, equal densities of each strain  $\{\text{Col} = 33\frac{1}{3}, \text{Resistant} = 33\frac{1}{3}, \text{Sensitive} = 33\frac{1}{3}\}$  is equivalent to the optimum RPS strategy of playing each move with equal probability, whilst a neighbourhood containing only Resistant bacteria  $\{0, 100, 0\}$  is equivalent to the strategy “always play Paper”. Figure 7.20 shows four CIAO plots of the *E. coli* simulation. The CIAO plot data is calculated by comparing the local density at each timestep against the local density at each previous timestep. To reduce the size of the CIAO plot, data is sampled every 50 timesteps (hence the plot is 100 pixels wide and deep). Each density comparison is evaluated as the expected result in a Rock–Paper–Scissors contest: results in favour of the later timestep ( $y$  axis) are dark. For example, if local density  $D_1 = \{100, 0, 0\}$  is compared with the *later* local density  $D_2 = \{50, 50, 0\}$ , the expected result will be 0.75 in favour of  $D_2$ ; thus giving pixel  $(D_1, D_2)$  a grey value 75% of maximum darkness.

From left to right, the CIAO plots display the raw data, Gaussian blur, zero-crossings, and fully processed image. The raw data plot (far left) exhibits both diagonal banding and tartan structures. The binarized image (far right) is much easier to interpret. The clear diagonal banding suggests that the system is exhibiting fairly regular cycling in this region of the lattice. The vertical “fault” line down the centre of the plot is an artefact of sampling: without sampling this feature does not exist. However, unsampled CIAO plots are too large to appear here.

The event plot of figure 7.21 highlights the points (sampled at a 50-timestep resolution) at which the local density of strain R, or strain C, exceeds 66%. The sequence “Resistance follows Col toxic” occurs regularly enough to again suggest fairly regular cycling in this region of the lattice. Figure 7.22 displays the local density of each strain over the first 1000 timesteps (plotted at a 50-timestep resolution); once again, regular cycling




Figure 7.21: Study 4 results. Event plot labelling the points (sampled at a 50-timestep resolution) at which the local density of Resistant (R) bacteria or Col toxic (C) bacteria exceeds 66% during one representative run. The repeated “R follows C” event sequence suggests regular cycling.




Figure 7.22: Study 4 results. Graph plotting the local density of each of the three bacterial strains during the first 1000 time steps (at a 50-timestep resolution). Cycling can be observed with regular sequence CRSCRS.

is clear. Finally, the state transition table of figure 7.23 also exhibits regular cycling.<sup>2</sup>

In summary, a suite of visualization techniques, including (for the first time) CIAO plots, have enabled the exploration of the nature of cycling suggested in the initial system snapshots of a biological model. In particular, the regularity and period of this cycling has been determined. The image processing techniques introduced in section 7.3.1.1 have been very useful in achieving this.

Given the discussion in section 7.3.4, of reasons why we would not expect regular cycling to be exhibited particularly often by coevolutionary systems, why is it that the model of *E. coli* coevolution generates such regular, periodic behaviour? There are two candidate explanations. First, the *E. coli* simulation employs a very simple representation of the space of possible strategies. Even the simple RPS system analysed in study 1 employed a discrete three-dimensional space of  $100^3$  possible strategies, whereas the *E. coli* system implements what is basically a simple set containing three strategies. There is no equivalent of the genotype space employed in our previous studies.

However, even if the *E. coli* model involved a more complex treatment of genotype space, coevolutionary cycles could remain regular, since the three bacterial morphs are

<sup>2</sup>Here, a state is entered when the local density of any one strain rises above 66% and exited once density falls below 66%. This classification scheme may be responsible for the 4 same-state transitions, without which cycling appears much more regular.

		To		
		R	P	S
From	R	2	9	0
	P	3	2	7
	S	5	1	0

Figure 7.23: Study 4 results. State transition table demonstrating regular cycling.

typically always present in the evolving population. While the strategy being played in each cell of the  $100 \times 100$  world varies over time, the global frequency of each strategy remains non-zero and roughly constant.

The update rules for the model ensure that if, for instance, all S individuals were to die, this strategy could never re-enter the population—only strategies that are already present may reproduce and spread to unoccupied cells since the model includes no analogue of mutation. Since, for the particular parameter values reported here, all three possible strategies are ever present, the system never spends time searching genotype space for adaptations and counter-adaptations; they are already present somewhere in the population. The cycling evident in the population time-series data reflects the time taken for a strategy to “migrate” across the grid, exploiting inferior competitors. At the relatively high population densities employed in the model, this leads to regular, periodic waves of succession.

The game-theoretic Evolutionarily Stable Strategy models of evolutionary dynamics typically employed in theoretical biology (Maynard Smith, 1982) also have no explicit representation of genotype space or mutation. In effect, such models assume that every possible strategy is always present at some non-zero (but perhaps infinitesimal) frequency within the population. All that varies in such models are the rates of reproduction enjoyed by these strategies. This ensures that the factors here identified as responsible for the irregular period of coevolutionary cycling (disengagement, neutrality, etc.) are not typically considered. The very different nature of the finite (co)evolving populations simulated in the adaptive behaviour literature offers an important opportunity to explore these factors.

## 7.5 Revisiting the Coin Toss Game

In this section, CIAO plots are used to visualize the results from section 6.2—the coin toss game—previously displayed in figure 6.6. The reason for this is two-fold. Firstly, CIAO plots may offer new insight into coin toss dynamics, affording better understanding of the effects of reduced virulence and resource sharing. Secondly, it will be interesting to observe CIAO plot visualizations of cycling, mediocre stability, engagement and funnelling. Offering a variety of dynamics, the coin toss data set is excellent for studying CIAO plot visualizations.

For all plots, “full” CIAOs are shown: ancestral elite (AE) contests across *all* generations are plotted, producing square rather than triangular output (see figure 7.24). To reduce the size of images, CIAO plots sample AE contests every 3 generations, producing  $200^2$  pixel plots. Each AE competition consists of 25 coin toss games between H-elite (host) and P-elite (parasite) opponents. Each game, the P-elite “tosses a coin” probabilistically, playing Heads or Tails according to the probability function 6.1 (displayed in figure 6.5). Results in favour of H-elites are output light, those in favour of P-elites are dark.

**Figure 7.24 (top, left): cycling with stasis.** Maximum virulence without resource sharing leads to cycling in the coin toss game (figure 6.6). However, this cycling is not regular; it is characterised by long periods of stasis—with parasites outperforming hosts—followed by rapid state transitions. For the majority of coevolution, the system is near disengagement. Unsurprisingly (see discussion in section 7.3.4), irregular cycling produces a tartan plot. The regularity of the plot results from the strict order of state transitions: “Heads to Tails”, “Tails to Heads”. Complementing figure 6.6, figure 7.24 nicely supports the inferences of section 6.2.3.

**Figure 7.24 (top, right): mediocre stability.** Maximum virulence with resource sharing leads to mediocre stability, or collusive disengagement, in the coin toss game (figure 6.6). Throughout coevolution, victories are shared between hosts and parasites. The resulting CIAO plot is very interesting. With the system at equilibrium throughout, one may expect AE contests to produce a homogenous grey CIAO plot. However, small fluctuations in elite strategies produces a fine-grained tartan pattern. Without the aid of figure 6.6, this plot is ambiguous, potentially showing stability, irregular cycling, or drift. Once again, the importance of using complementary visualizations to support CIAO plot data is confirmed.

**Figure 7.24 (bottom, left): engagement.** In the coin toss game, reduced virulence without resource sharing results in engagement in the most difficult (unpredictable) part of

space (figure 6.6). When competing with recent ancestral P-elites, H-elites perform well; indicated by the white banding observed on the leading diagonal. Throughout coevolution, hosts are always capable of competing with parasites. Continually striving to “escape” hosts, parasites relentlessly move through difficult regions of genotype space. As hosts chase, previous adaptations are lost. This is indicated by the prevalence of dark shading throughout the plot: if adaptations were preserved, the plot would be much lighter. Whilst hosts can compete with current parasites, they are not good at predicting parasites from much earlier, or future, generations: there is short term evolutionary memory. This additional insight is not available from figure 6.6.

**Figure 7.24 (bottom, right): funnelling.** Finally, reduced virulence with resource sharing produces funnelling: after a period of mediocre stability, populations engage in the most difficult part of space (figure 6.6). Around generation 440, a change in H-elite strategy produces a dramatic transition from mediocre stability to engagement. Before now, such transitions have not appeared in the CIAO plot literature. As such, in isolation, the plot is difficult to interpret; however, when considering the suite of plots, in conjunction with the alternative visualizations of figure 6.6, interpretation becomes clear. As an example of funnelling, this plot may prove a useful reference for future CIAO plot studies.

**Conclusions:** CIAO plots offer expressive and informative visualization of coevolutionary dynamics. However (reiterating the conclusions already reached in this chapter), to be interpreted correctly and unambiguously, it is necessary to complement their use with other (perhaps domain specific) visualizations.

## 7.6 Steering Coevolution Through Interactive Visualization

Typically, evolutionary algorithms are run in batch mode: once initial conditions have been decided, the evolutionary process is left to run without further human intervention (for example, see the simple GA of section 2.2.1). Upon termination summary statistics are presented, often in the form of simple graphs, and often disguising much critical (mid-run) information. Significant interaction with the algorithm can only take place at this stage: on the basis of crude and fragmented end-of-run visualizations, starting parameters (initial conditions) may be varied for subsequent runs. This method is employed throughout chapters 6 and 7. Unfortunately, this laborious process is not guaranteed to find the required parameter settings.

At the opposite extreme, some algorithms require user interaction at every stage. Each

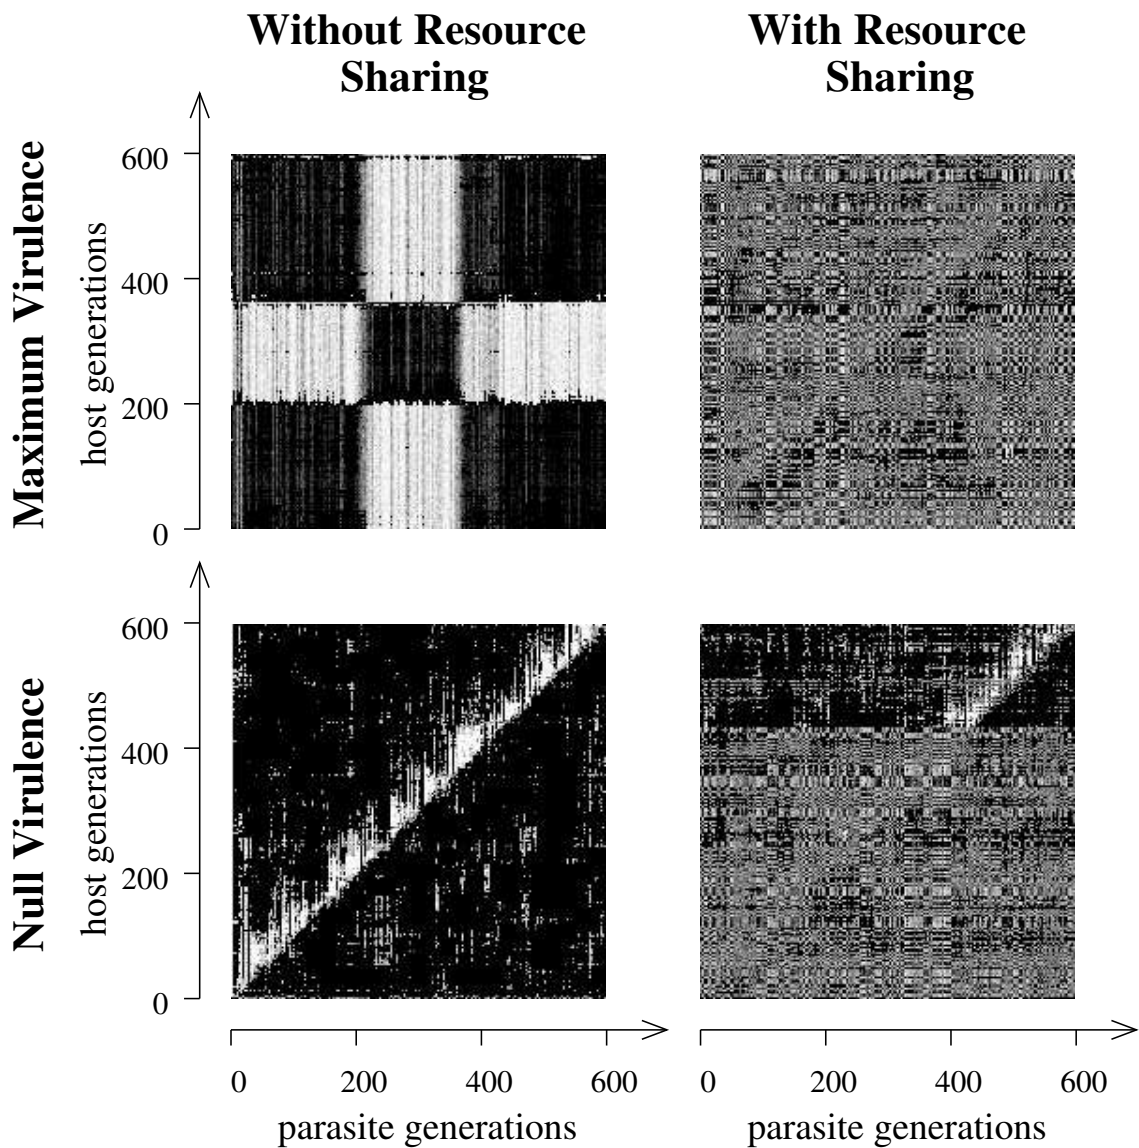


Figure 7.24: CIAO plots visualizing typical coevolution of the coin toss game. Top left: irregular cycling with periods of stasis and sharp transitions. Top right: mediocre stability. Bottom left: hosts engage parasites in difficult regions of space. Bottom right: a transition (funnelling) from mediocre stability to engagement occurs around host generation 440.

generation users select which members of the population will contribute genetic material to the following generation. Without user interaction, the system halts. Hence, these algorithms demand much human input. As such, they are typically only considered when an accurate fitness function is difficult, or impossible, to operationally define. Examples include aesthetic evolution evolved to users' taste (Dawkins, 1986), computer graphics (Sims, 1991) and works of art (Dorin, 2001).

Computational steering—here defined as the ongoing, informed intervention of users in the execution of an otherwise autonomous computational process (see, e.g., Parker, Johnson, & Beazley, 1997)—lies somewhere between these two extremes. Algorithms may be guided by human intervention, when appropriate, but will otherwise continue autonomously; hence, the degree of intervention—altering system parameters, or turning on or off aspects of the algorithm, etc.—is decided by the user. As such, computational steering offers the power of intervention without the demands of forced interaction. Computational steering tends to be used in situations where the skills and knowledge of human users are critical to system performance, but are difficult (or impossible) to operationalise. If it *were* possible, it may not be necessary to rely on humans in the system at all.

### 7.6.1 Prospects for Computational Steering

Computational steering of (co)evolutionary algorithms is potentially beneficial for several reasons. Foremost, expert guidance through human intervention may well improve the quality of evolutionary algorithms for design and problem solving; as knowledge uncaptured in the algorithm itself is delivered by hand. In conjunction, steering may lead to insights into how systems work; as experts perform run time manipulation, tacit knowledge may be teased out. As a secondary effect, a drive towards evolutionary steering is likely to encourage the development of novel visualization techniques that are beneficial in their own right. Finally, steering devices offer an excellent method of teaching novices about systems; allowing students to gain practical experience of evolutionary phenomenon in an engaging and intuitive way.

However, to make evolutionary steering work, it is critical that users are able to make *effective* interventions. To achieve this, three conditions are necessary. Firstly, users must have the *expertise* to effectively steer computation: to guide evolution appropriately, users require a strong theoretical grounding. Secondly, users must be able to exert the required influence in a suitable manner: if the necessary means of *interaction* are missing, expert users become redundant; for example, knowing that mutation rate should be lowered is little use if such functionality is not available. Finally, users must be presented with the



information they require in a similarly intuitive manner: *visualization* must offer users the stimulus they need to impart necessary guidance (for a more detailed discussion of these principles, and treatment of the prospects of evolutionary steering in general, refer to Bullock, Cartlidge, & Thompson, 2002).

In the following section, a first attempt at coevolutionary steering is presented. This simple  $\alpha$ -model demonstrates an example interface for steering the virulence of populations coevolving in the numbers game (refer to section 6.1).

## 7.6.2 Proof of Concept

### 7.6.2.1 Steering the Numbers Game

Here, a novel coevolutionary steering tool is presented. Utilising the simple numbers game of section 6.1, an interactive user interface is developed, enabling users to vary the virulence of parasites (and hosts) and other evolutionary parameters during run time. Inherently simple and easy to visualize, the numbers game is an ideal domain to demonstrate coevolutionary steering. Further, in the drive for dynamic virulence (section 6.4), steering potentially offers an alternative solution to autonomous fitness variation; if not directly, then at least insofar as steering may help tease tacit knowledge out of expert users, providing a greater theoretical framework on which to develop future dynamic virulence techniques.

Figure 7.25 shows the graphical user interface designed to steer the coevolutionary numbers game. **A:** the virulence curve of parasites, showing the relationship between parasite score and parasite fitness (refer to figure 5.2); also used to display the distribution of parasite scores (see figure 7.26). **B:** the virulence curve of hosts used to show the relationship between host score and fitness. Unlike previous models, the fitness of hosts is also subject to virulence (the biological analogy breaks down here somewhat). **C:** slider used to vary parasite virulence,  $\lambda_p$ , between 0.5 and 1.0 (Null to Maximum), displayed in graph A. Here,  $\lambda_p = 0.59$ . **D:** slider used to vary host virulence,  $\lambda_h$ , displayed in graph B. Here,  $\lambda_h = 0.94$ . **E:** slider used to vary parasite bias,  $B_{par}$  between 0.0 and 1.0. Host bias cannot be varied and is *always* 0.5. **F:** slider to vary the speed of coevolution. When run at slow speeds, the interface affords more user control. **G:** the set of initial parameters: once a run begins, these parameters are fixed. Runs may be repeated by reusing the same random seed. Random seeds may be input manually, or selected at random using the “random” button. When ticked, the scaling box normalizes the scores of hosts and parasites with respect to the highest scoring member of each population before utilising the virulence function (as performed throughout this thesis). Without scaling, scores are

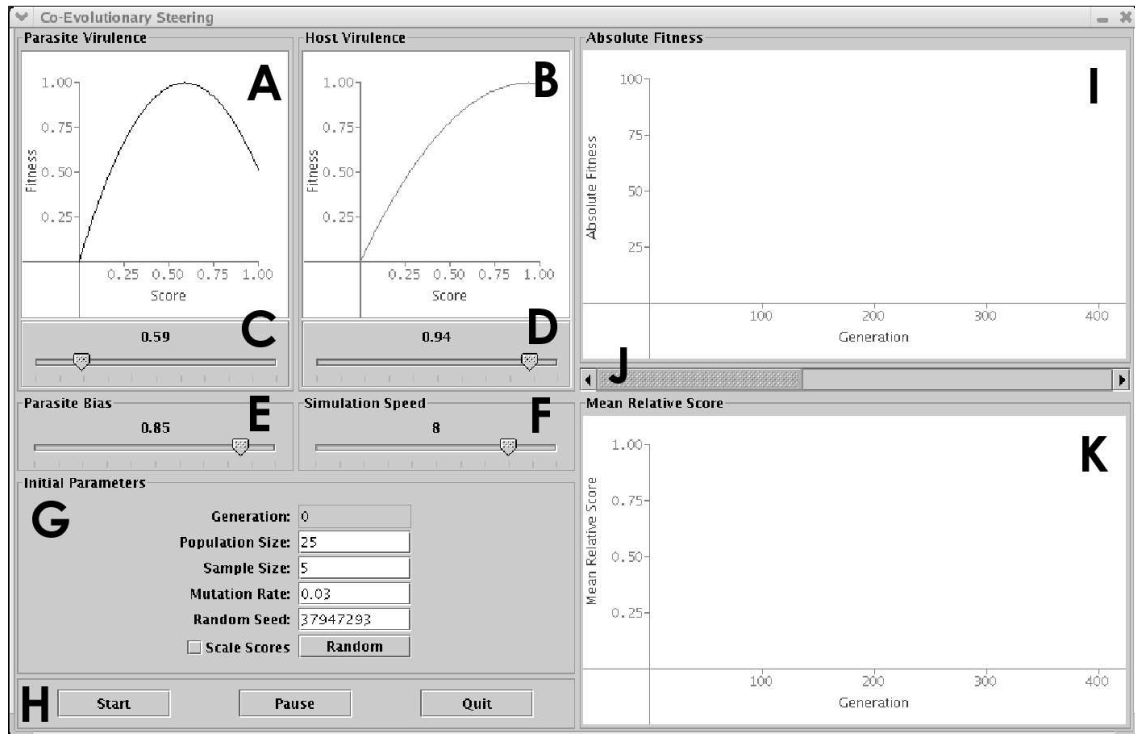


Figure 7.25: Coevolutionary steering interface.

input directly. **H**: control buttons to start a run, pause/continue, or exit the interface. **I**: absolute fitness graph displaying the  $n$  values of each individual; equivalent to the large (top) graphs of figures 6.1 and 6.2. **J**: scroll bar allowing users to scroll across generations of graphs I and K. **K**: relative fitness graph displaying the mean relative score of each population; equivalent to the small (bottom) graphs of figures 6.1 and 6.2. In all figures, parasites are depicted in dark shade, with hosts in light.

The virulence graphs of the steering interface (figure 7.25, A and B) are also used to visualize the distribution of relative fitness scores in each coevolving population. This is achieved by effectively overlaying a histogram of fitness scores on top of the virulence curves: at each fitness score, a bar is drawn in proportion to the number of individuals attaining that fitness. Figure 7.26 shows two schematical examples. In each graph, a Null virulence curve ( $\lambda = 0.5$ ) is shown, however, the value of  $\lambda$  does not affect the size of the distribution bars, only their position: each bar is always centred on the curve. The graph to the left shows a uniform distribution with each bar of equal size. This shows that equal numbers of individuals are achieving each of the possible relative fitness scores (between 0.0 and 1.0). To the right, an exponential distribution is shown. Here, far fewer individuals score high fitness values than those that score low: as relative fitness increases, the size of the distribution bars gets smaller; most individuals score 0.0 (indicated by large




Figure 7.26: Example schematics of relative fitness distributions. Left: uniform distribution of fitness scores; right: exponential distribution. Bars are centred on the virulence curve, with length proportional to the number of individuals attaining each fitness score. These lengths are independent of virulence; shown as  $\lambda = 0.5$  in both graphs.

bars) and none score 1.0 (no bar). This novel visualization method is intended to present users with information regarding the underlying distribution of fitness scores within each population; potentially useful when determining the appropriate level of virulence.

### 7.6.2.2 Example Output

This section describes an example coevolutionary run utilising the interactive steering interface previously described. The purpose of this section is to demonstrate the functionality of the steering interface. Unlike previous sections, there is no attempt to either coevolve successful individuals, or to maintain engagement. The example is intended only to give the reader an insight into relationships between user interactions and system behaviour.

**Initial conditions.** At the beginning of the run, the following initial conditions are input into the interface: population size 25, sample size 5, mutation rate 0.03, random seed 10000 and “scale scores” on. Once the run begins, these parameters may *not* be altered.

The following variable parameters are chosen to begin: parasite virulence  $\lambda_p = 0.75$ , host virulence  $\lambda_h = 0.75$ , and parasite bias  $B_{par} = 0.9$ . For the first 100 generations, these values are unaltered.

**Screen shot 1: figure 7.27** shows the system at generation 100. The graph of absolute fitness (top, right) shows both populations quickly evolve to approximately  $n = 85$  1-alleles. The populations are engaged throughout—shown in both the absolute and relative fitness (bottom right) graphs—with parasites (dark shade) having the upper hand; a result of the large asymmetrical bias,  $B_{par} = 0.9$ . The relative fitness distributions of both populations

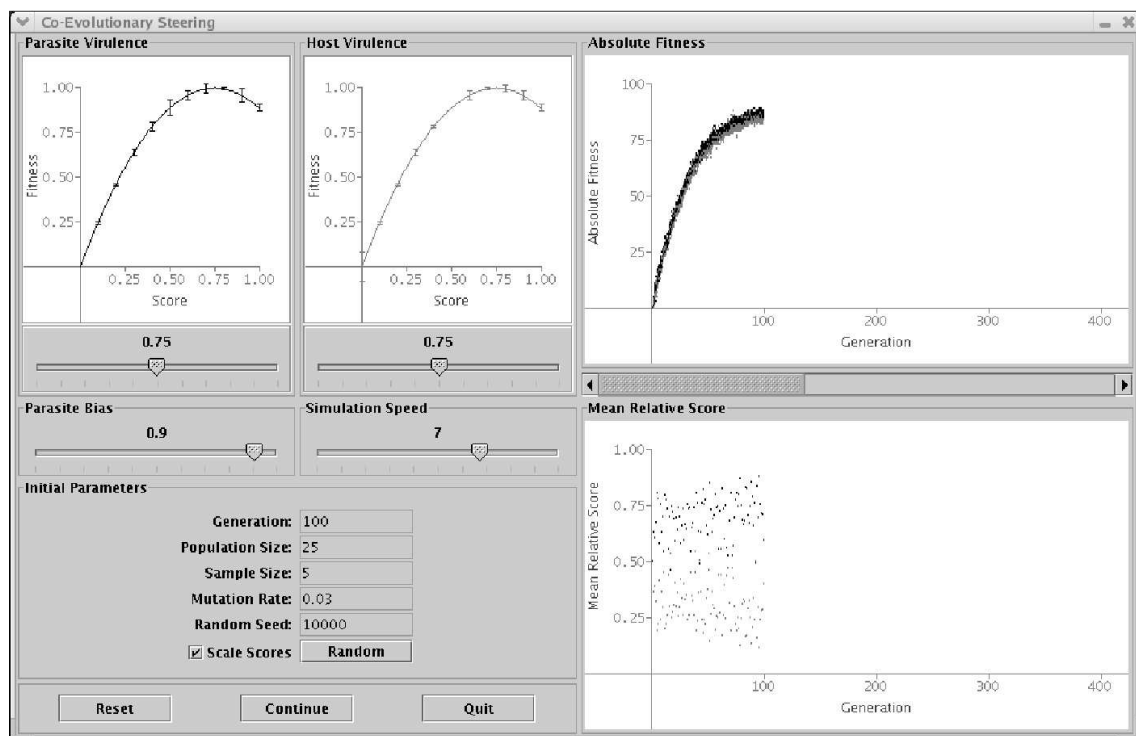


Figure 7.27: Screen shot 1. Generation 100.

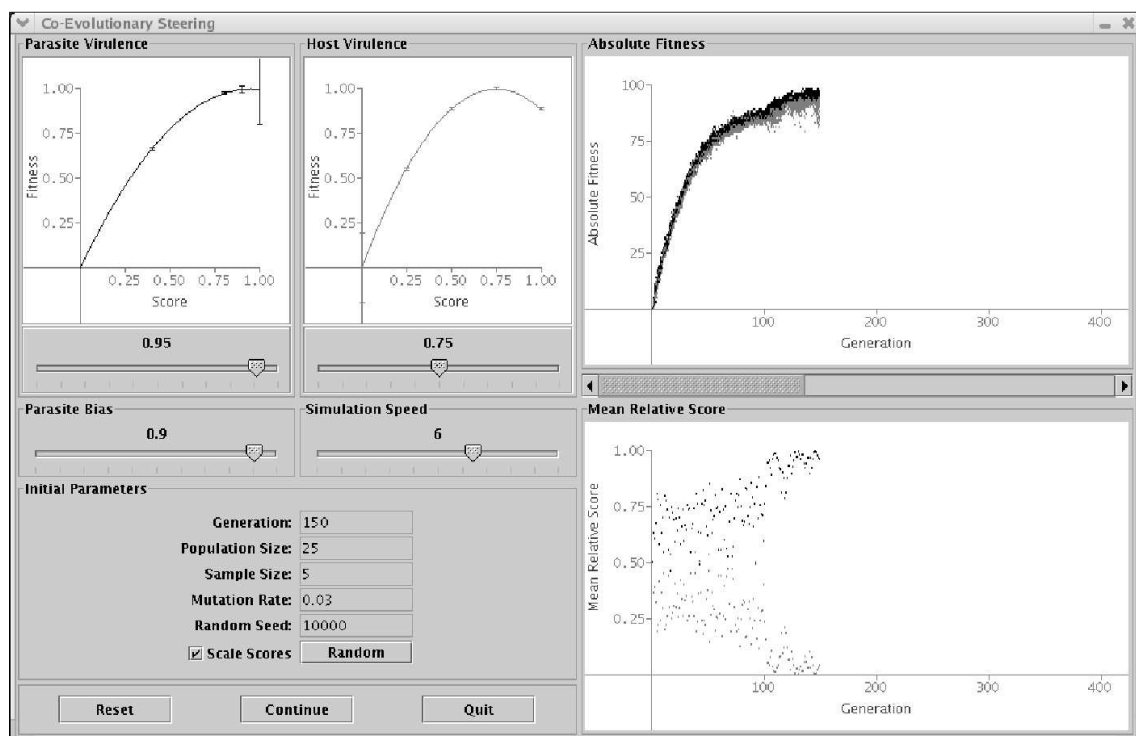


Figure 7.28: Screen shot 2. Generation 150.

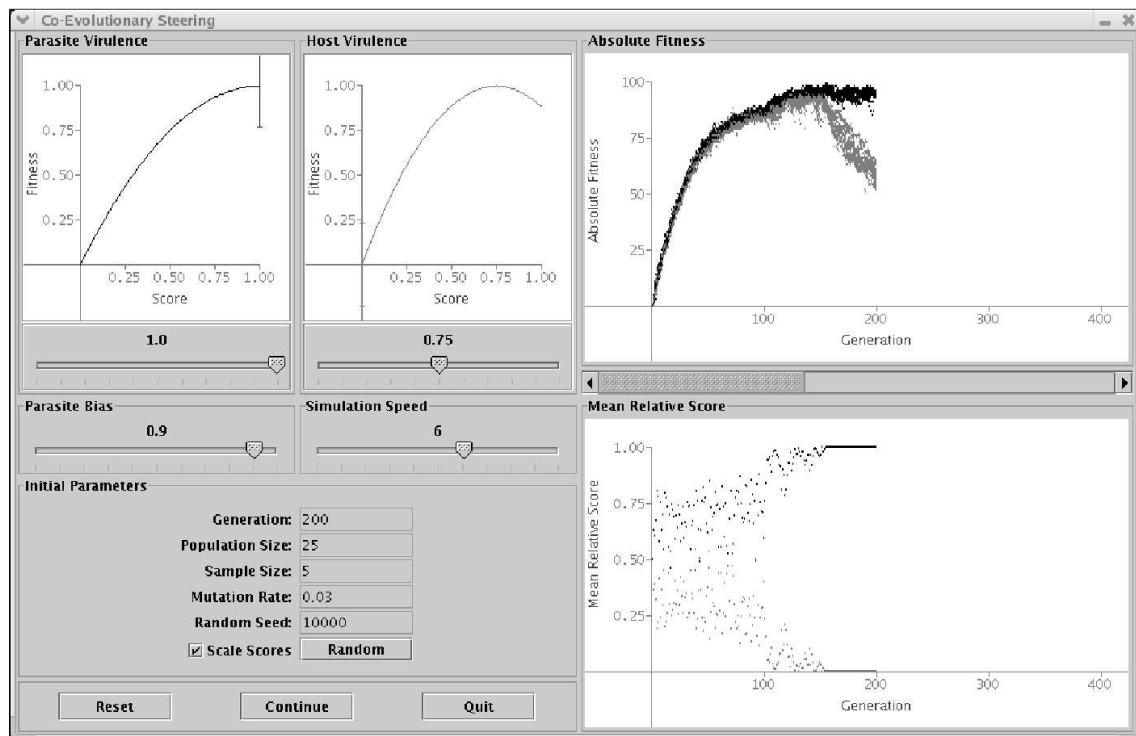


Figure 7.29: Screen shot 3. Generation 200.

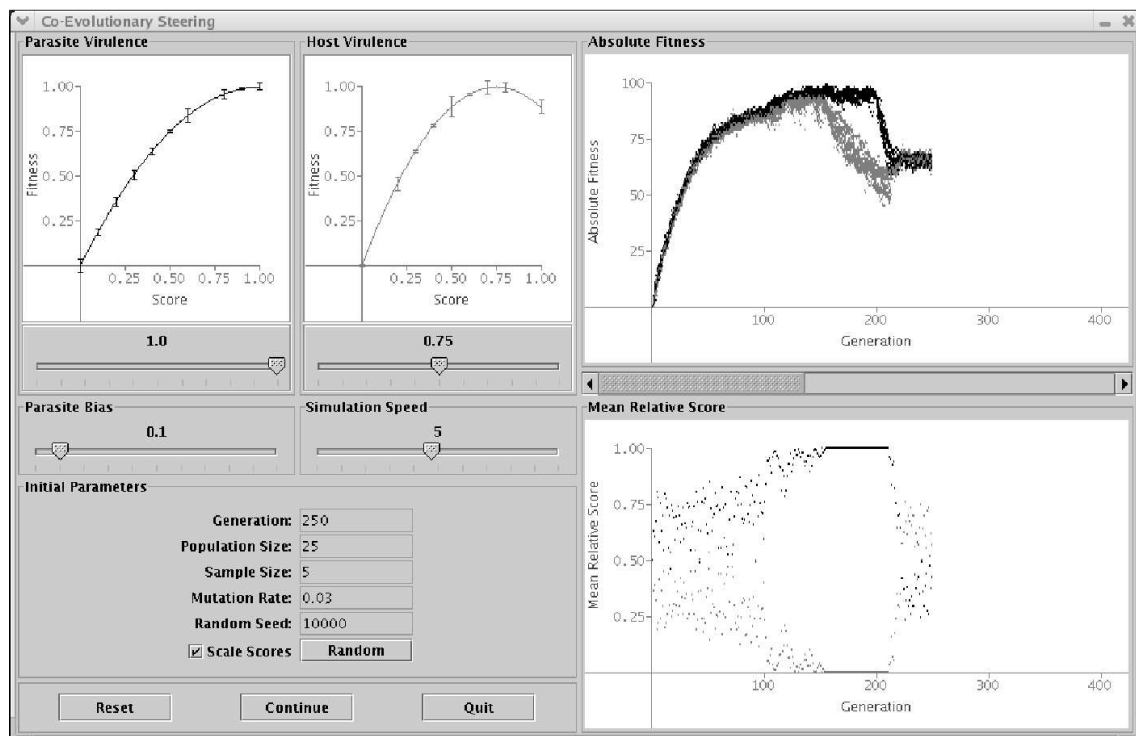


Figure 7.30: Screen shot 4. Generation 250.

is fairly uniform (top; left and centre); there appears little danger of disengagement. As a result, parasite virulence is increased to  $\lambda_p = 0.95$ .

**Screen shot 2: figure 7.28** shows the system at generation 150. The increase in  $\lambda_p$  has increased the  $n$ -values of both populations (top, right), however, as a consequence, the populations are nearing disengagement. This can be observed in the graph of mean relative fitness (bottom right) and the distribution graphs: nearly all parasites score 1.0 (top, left) whilst the majority of hosts score 0.0 (top, centre). However, parasite virulence is further increased to Maximum;  $\lambda_p = 1.0$ .

**Screen shot 3: figure 7.29** shows the system at generation 200. The populations have disengaged (observable in the distribution and relative fitness graphs). This occurred soon after parasite virulence was increased to Maximum. As a result, the absolute fitness of both populations has fallen towards the mutational baseline levels of  $n = 50$  for hosts,  $n = 90$  for parasites. In response, parasite bias is dramatically lowered to  $B_{par} = 0.1$ .

**Screen shot 4: figure 7.30** shows the system at generation 250. The reduction in  $B_{par}$  at generation 200 has shifted asymmetrical advantage in favour of hosts. Initially, the absolute  $n$ -values of parasites fall towards the baseline mutation bias, however, upon reengaging with hosts, both populations progress to around  $n = 65$ . To increase the system's  $n$ -values, host virulence is increased to Maximum;  $\lambda_h = 1.0$ .

**Screen shot 5: figure 7.31** shows the system at generation 275. The increase in  $\lambda_p$  has pushed the system to the brink of disengagement.

**Screen shot 6: figure 7.32** shows the system at generation 350. The system is disengaged. Both populations have reached their baseline levels of  $n$ . The system is unlikely to reengage under current conditions. As such, parasite bias is increased to  $B_{par} = 0.8$ ; returning asymmetrical advantage back to parasites.

**Screen shot 7: figure 7.33** shows the system at generation 400. After increasing  $\lambda_p$ , the populations have reengaged and the system is performing well, averaging  $n$ -values around 85. In an attempt to maximise  $n$ , parasite bias is increased to  $B_{par} = 0.95$ . To compensate, parasite virulence is marginally decreased to  $\lambda_p = 0.95$ .

**Screen shot 8: figure 7.34** shows the system at generation 500. The system is engaged at near maximum values of  $n$ .

### 7.6.2.3 Future Work

The demonstration platform presented in this section has been used to highlight the potential utility of (co)evolutionary steering. However, this  $\alpha$ -model is very much the beginning of this line of study. The majority of work lies ahead.

In future, this platform will be used to perform several studies. Firstly, the interface

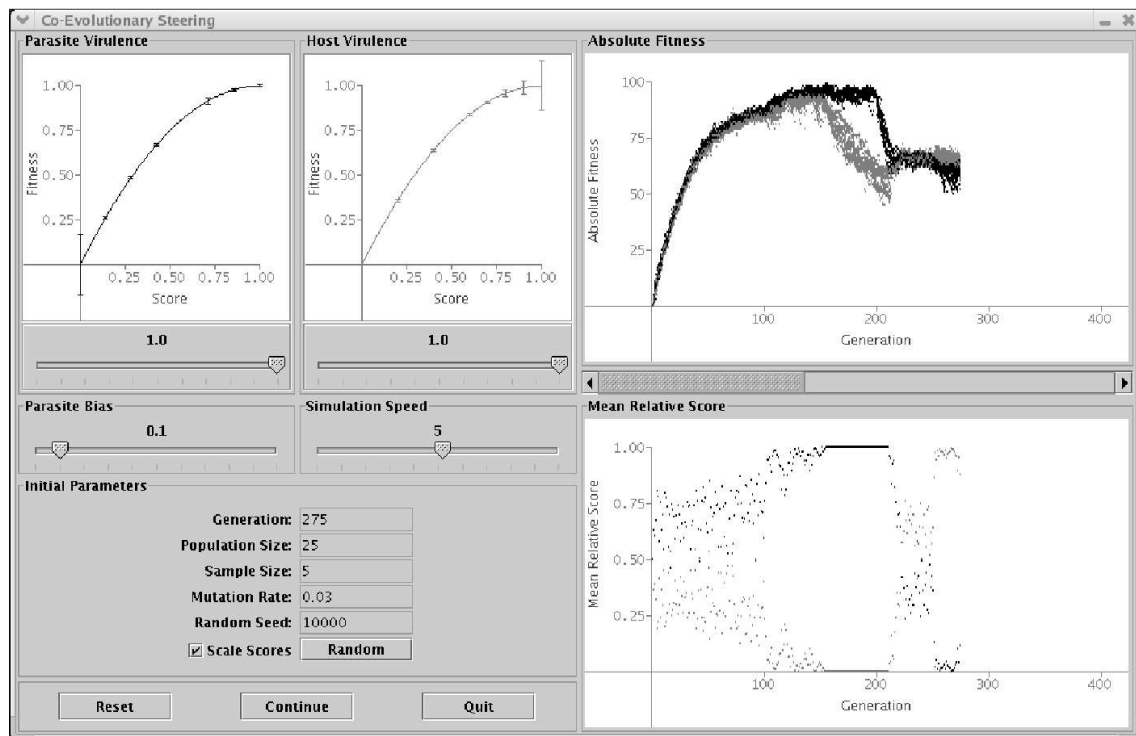


Figure 7.31: Screen shot 5. Generation 275.

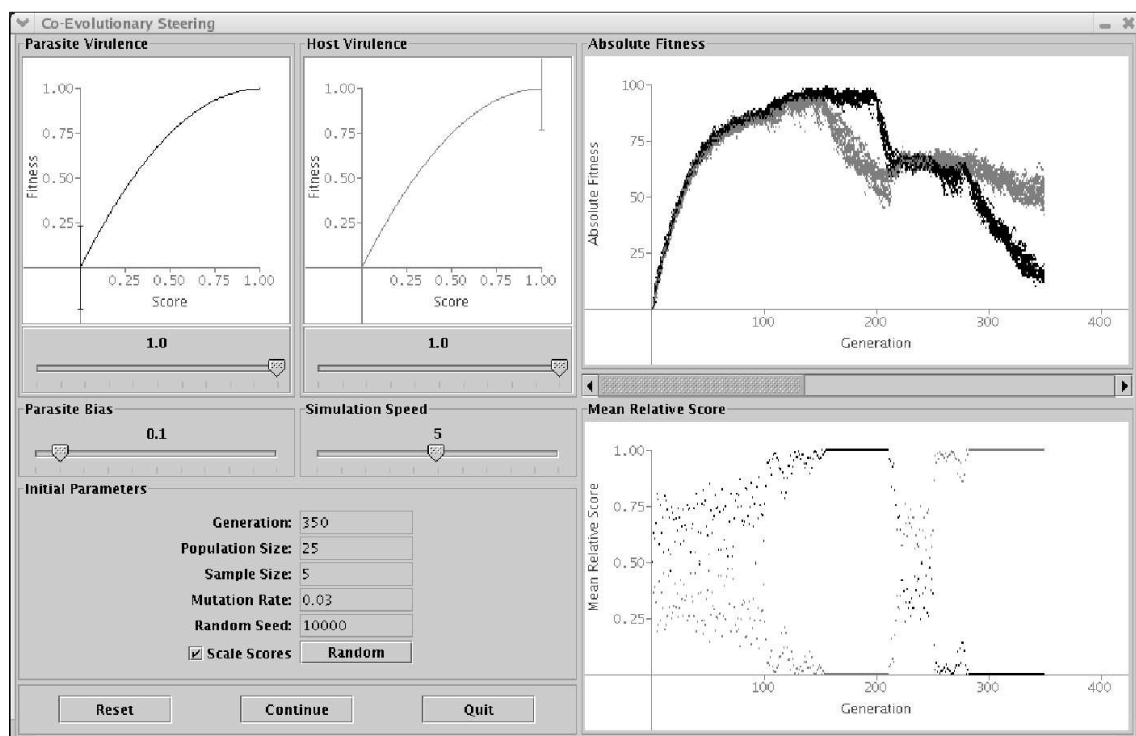


Figure 7.32: Screen shot 6. Generation 350.

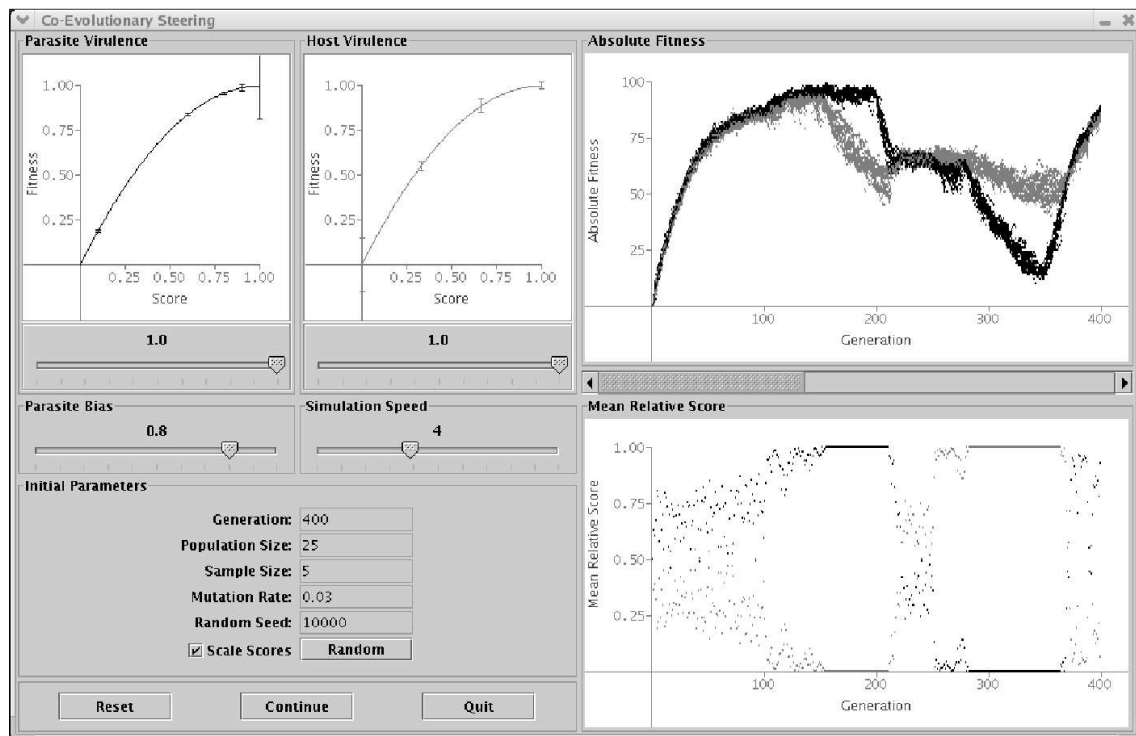


Figure 7.33: Screen shot 7. Generation 400.

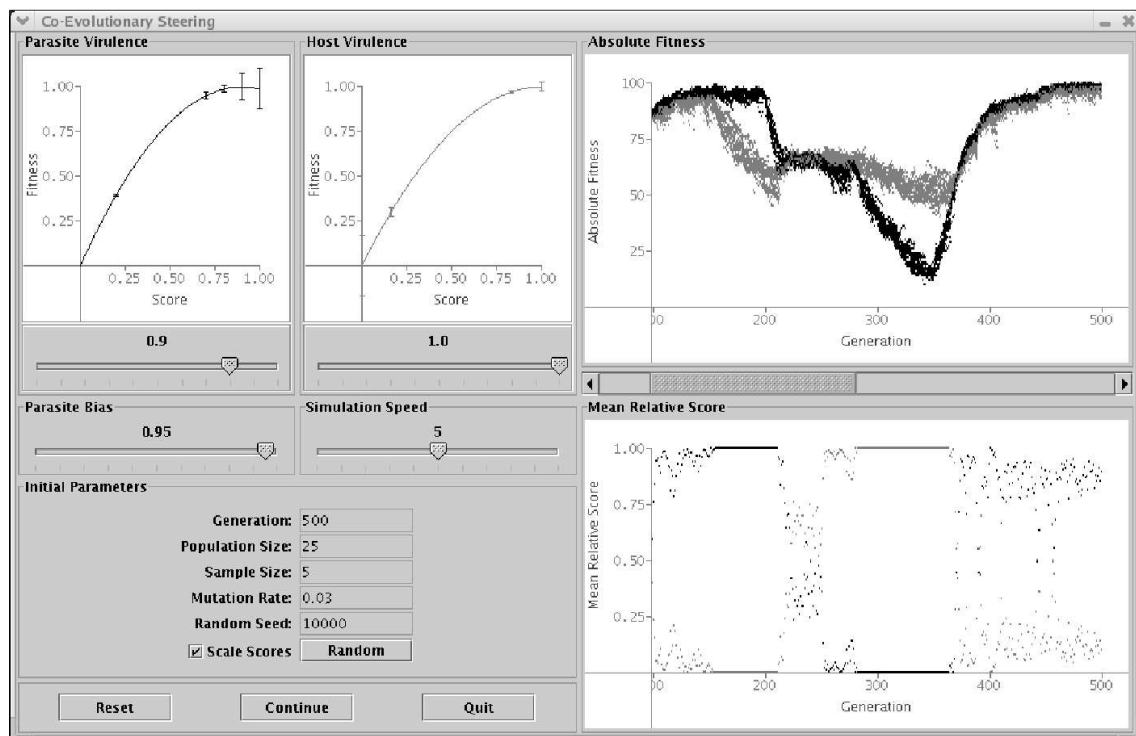


Figure 7.34: Screen shot 8. Generation 500.



will be modularised, allowing parts of the interface to be removed, or fixed, thus enabling more or less complex interfaces to be easily constructed. Experiments will compare the performance (maximum stable  $n$ -value achieved) of naive users with experts as they attempt to effectively steer coevolutionary under various conditions. These subjects will be carefully debriefed in an effort to understand how they went about their task. Any insights gained could be used to design more efficient automatic coevolutionary algorithms (i.e., dynamic virulence; section 6.4).

In addition, the platform will be used directly as a tool with which to explore coevolutionary disengagement. The nature of the platform allows a rich interaction with the coevolutionary system, which will hopefully lead to improved insights into algorithm behaviour. Finally, the steering platform will enable the teaching of coevolutionary concepts to students in a hands-on manner that will hopefully be engaging and informative.

This example of computational steering is very crude in many respects. The visualizations used to inform the user of algorithm behaviour are, for the most part, standard plots of fitness over time. In general, much more sophisticated indicators of a wider range of system aspects could be employed. Similarly, the range of interaction supported by the example steering platform is limited to varying a small number of the parameters governing evolutionary selection pressures. There are, of course, many alternative aspects of algorithm performance that could be influenced by the user. However, the simple steering interface presented here is adequate for the specific purpose of exploring the relationship between virulence and engagement in the minimal coevolutionary numbers game.

## 7.7 Summary

If coevolutionary algorithms are to ever reach their full potential, appropriate (effective) visualization techniques are necessary. However, visualization tools specifically designed to further the theoretical understanding of coevolutionary algorithms are rare. CIAO plots—primarily designed to highlight cycling—are an exception.

Cycling occurs in many adaptive systems. Mutational stochasticity, neutrality, disengagement and rugged fitness landscapes each contribute to the irregularity of these cycles. As a result, irregular cycling may characterise the adaptive behaviour of many unstable coevolutionary systems and may contribute to the failure of coevolutionary optimisation. Detecting and characterising this type of dynamic is difficult within modelling paradigms that privilege stable states (e.g., evolutionary game theory). As a result, coevolutionary adaptive behaviour simulation and visualization are good candidates for improving our understanding of these types of irregular cycling.

In this chapter, it has been demonstrated that CIAO plots, a widely used tool for visualizing coevolutionary progress, are difficult to interpret with respect to irregular coevolutionary cycling. Further, there are reasons to believe that irregular coevolutionary cycling is a common and significant category of coevolutionary dynamics. Hence, CIAO plots should preferably be used in conjunction with other visualizations: problem-specific analysis methods can usefully complement CIAO plots and can aid in their interpretation.

For the first time, CIAO plots have been applied to biological simulation data, affording valuable insights into the underlying coevolutionary dynamics that agree with empirical observations. This exemplifies the potential cross-over between techniques designed for evolutionary computation and adaptive behaviour more generally.

Computational steering offers a method of ongoing, informed intervention of users in the execution of an otherwise autonomous computational process. In this chapter an example steering interface is presented, allowing users to guide coevolution through *interactive* visualization. Allowing virulence to be manually manipulated during run time, coevolutionary steering offers an alternative approach to autonomously varying virulence functions, and may offer new insight into the relationship between virulence and engagement. To achieve this, however, effective visualizations are critical.

Working towards computational steering for evolutionary computation has considerable merit if there is some chance of achieving the improvements in algorithm performance and insights into algorithm behaviour that it may bring. However, even if such a research effort were to fail in its stated primary aim, in pursuing it, several secondary goals would be progressed.

First, effective visualization and interaction are worthy targets in their own right; as demonstrated by the investigation of CIAO plots. Second, and perhaps more importantly, in integrating visualization, interaction and EC theory, research into computational steering for evolutionary computation has the potential to improve our understanding of all three. To the extent that visualization and interaction techniques are motivated by some theory of evolutionary computation, the ability of these techniques to support computational steering is a good indicator of the success of EC theory in guiding our understanding of evolutionary algorithms.

# Chapter 8

## Transient Neutrality

---

Throughout this thesis, it has repeatedly been described how coevolutionary disengagement results in periods of evolutionary *drift*. Recently, within the evolutionary computation community, there has been some interest in the role of neutral drift in the behaviour of evolutionary algorithms (Barnett, 1998). It has been suggested that neutral drift might alleviate problems of premature convergence, and that proper appreciation of such drift necessitates a radical reappraisal of our picture of how evolutionary algorithms work (Barnett, 1998; Shipman, Shackleton, Ebner, & Watson, 2000; Smith, Husbands, Layzell, & O’Shea, 2002). How does the drift that results from coevolutionary disengagement relate to that experienced by populations *percolating* across “neutral networks”?

Within the evolutionary computation community, search-space neutrality has been defined as the property of adjacent points in a search space exhibiting equivalent fitness scores. A *neutral network* is a set of such points, where each member of the set neighbours at least one other member. Clearly, these notions of “adjacency” and “neighbourhood” must be understood in terms of a search algorithm’s genetic operators—often this will be a complicated matter. Moreover, the idea of “equivalent fitness” that lies at the heart of the neutrality concept is perhaps not as straightforward as it might appear (see, e.g., Bullock, 2002).

For certain real-world search spaces, such as the RNA folding map, it has been demonstrated that the neutrality present is of a potentially useful kind (Huynen, 1996; Huynen et al., 1996; Fontana & Schuster, 1998a, 1998b). Neutrality in this case stems from the fact that many RNA sequences (genotypes) fold into the same secondary structure (phe-

notypes). This redundancy ensures that some mutations may alter a genotype without altering the associated phenotype. It just so happens that neutral networks *percolate* the RNA space, ensuring that a large proportion of possible phenotypes (RNA secondary structures) are accessible from an arbitrary genotype via drift. Furthermore, RNA neutral networks enjoy a property of *constant innovation* in that, over many generations, a neutral walk across such a network will tend to encounter novel phenotypes at a constant rate comparable to that which would be achieved by a random walk in the search space. These properties would seem to ameliorate RNA evolution. Could similar properties be exploited by evolutionary optimisation algorithms?

Notice that in this canonical example of neutral drift, the notion of “equivalent fitness” is understood in terms of “equivalent phenotype”. Two adjacent RNA strands that fold into the same structure are clearly part of a neutral network, since they code for phenotypes that are practically identical. However, since different phenotypes may nevertheless achieve the same fitness score, it is quite possible that a single neutral network may include many *different* phenotypes that are “selectively neutral” with respect to one another. Although this complication is acknowledged within the neutrality literature, it is not given much attention.

In evolutionary (rather than coevolutionary) systems, if two phenotypes achieve the same fitness score, one has reasonable grounds for classifying them as “equivalent” in some important and enduring sense—selection (ignoring noise and sampling error) will never discriminate between them—neutrality is *inherent*. However, in a coevolutionary scenario, the fact that two phenotypes achieve equivalent fitness is less compelling—since fitness intrinsically depends upon opponents faced, it is entirely possible that they may never achieve equivalent fitness again—neutrality is *transient*.

In fact, where fitness is calculated relative to performance against a coevolving population, the notion of neutral drift requires a significant overhaul. No longer can one consider neutral networks to be a (fixed) property of a fitness landscape. Rather, neutrality is now a relational property predicated on the current makeup of both populations. Neutral networks are transient phenomena, merely reflecting the subjective (and hence temporary) equivalences that may exist between the current conspecifics *with respect to their current coevolutionary opponents*. As a result, whereas neutral drift in an *evolutionary* system is likely to preserve the quality of evolved solutions by restricting genotypic change such that phenotypes are equivalent in some “objective” sense (e.g., they code for the very same physical structures), the same drift in a coevolutionary system is not necessarily so constrained. In particular, the drift caused by coevolutionary disengagement may involve genotypic changes that correspond to radical phenotypic change, just so long

as these phenotypic changes are “undetected” or “invisible” to the individuals that happen to make up the contemporary coevolutionary partner population.

In many cases, given that a disengaged population will often be more likely to generate more disengaged offspring than offspring that re-engage with their coevolutionary partner population, this type of *disengaged drift* may resemble random movement through the local genetic neighbourhood of the disengaged population. Given that, in the period prior to disengagement, a population contains reasonably high-quality individuals, far from preserving phenotypic quality, this type of disengaged drift is likely to be deleterious.

The exact character of any evolutionary drift is likely to be influenced by the nature of the genetic encodings and genetic operators employed (Bullock, 1999, 2001). Since these encodings and operators may be different for each population involved in a coevolutionary algorithm there is no guarantee that disengaged drift will readily lead to re-engagement. It is possible that encoding/operator biases ensure that drifting populations tend to remain within portions of genotype space that effectively maintain disengagement (as was observed, for example, in the numbers game, section 6.1). In general, the issues briefly raised here suggest that our appreciation of neutrality and drift and how these phenomena apply to coevolutionary systems deserves more careful consideration. In the following sections, a potential framework for studying these (coevolutionary) phenomena is considered.

## 8.1 Modelling Neutrality With *NK* Landscapes

Providing a tuneably rugged family of model fitness landscapes, the *NK* framework was developed to investigate the effects of landscape ruggedness on evolving populations (Kauffman, 1993). In Kauffman’s *NK* model,  $N$  corresponds to the number of genes in a genotype (or traits in an organism), while  $K$  corresponds to the number of other genes (traits) which affect the fitness contribution of each gene (trait). Thus,  $K$  corresponds to the richness of epistatic linkages. If  $K$  is increased from low to high, the *ruggedness* of the landscape increases: i.e., the number of fitness peaks increases, the steepness of the sides of fitness peaks increases and the mean height of fitness peaks decreases. By altering  $N$  and  $K$ , the *NK* framework provides a family of landscapes which range from the single peaked and highly correlated ( $K = 0$ ) to the fully random and uncorrelated ( $K = N - 1$ ).

Since the fitness contribution of each trait depends upon itself and epistatically on  $K$  other traits, there are  $2^{K+1}$  traits which bear upon its fitness. Thus, each trait’s fitness contribution is determined by the particular combination of “inputs”. For each of the possible

$2^{k+1}$  input combinations to the  $i$ th trait, a real valued fitness contribution is assigned in the range  $0 \leq w_i < 1$ ; to capture the complexity of epistatic interactions, contributions are assigned randomly. The fitness of the entire organism (genotype) is calculated as the mean contribution of each of the  $N$  traits (loci); also in the range  $0 \leq W < 1$ .

Figure 8.1, left, demonstrates an example  $NK$  landscape with  $N = 3$  and  $K = 2$ . As the diagram shows, since  $K = N - 1$ , each of the  $N$  traits are epistatically linked to *all* other traits. The fitness contribution of each trait,  $w_i$ , is assigned a real value at random. The fitness of the entire organism,  $W$ , is calculated as the mean value of each trait's fitness contribution. Notice that each  $W$  is different: there is no neutrality. Since fitness contributions are (randomly) assigned real values, the probability of any two being equal is extremely small. Thus, any single mutation will likely result in a different fitness contribution: there are unlikely to be any neutral mutations in an  $NK$  landscape. Before  $NK$  landscapes can be used to investigate neutrality, therefore, they require structural alteration.

Several models have been proposed to generalize Kauffman's  $NK$  framework so that a tuneable level of neutrality can be added. Perhaps the most pervasive are the  $NKp$  ("probabilistic"  $NK$ ) model, developed by Barnett (1998) and another (the  $NKq$ , or "quantised"  $NK$ ; as termed by Geard, Wiles, Hallinan, Tonkes, & Skellett, 2002)—developed by Newman and Engelhardt (1998). While both extend Kauffman's  $NK$  model, each employs a distinctly different approach to modelling neutrality.

The  $NKp$  model introduces a novel parameter,  $p$ , to the  $NK$  framework. This parameter controls the probability that a certain input combination will make *no contribution* to an organism's fitness. Thus, if  $p$  is high, many input combinations make no contribution. When a non-contributory input combination is mutated, mutation is neutral when the new input combination *also* makes no fitness contribution. The likelihood of neutral mutation—i.e., the amount of neutrality—increases with  $p$ . Figure 8.1, centre, demonstrates an  $NKp$  landscape with  $N = 3$ ,  $K = 2$  and  $p = 0.75$ . This landscape is identical to the  $NK$  landscape displayed to the left, except that the fitness contribution of 75% of input combinations have been reduced to 0.0. As a result, half of all possible organisms are neutral (have equal fitness,  $W$ ).

The  $NKq$  model introduces neutrality to the  $NK$  framework through the novel parameter,  $q$ . Rather than assign fitness contributions as random real numbers in the range  $0 \leq w_i < 1$ , Newman and Engelhardt (1998) assign random *integers* in the range  $0 \leq w_i < q$ ; thus if  $q = 2$ , say, each contribution  $w_i$  is either zero or one. The fitness  $W$  of the entire organism (genotype) is proportional to the sum of the contributions at each locus:

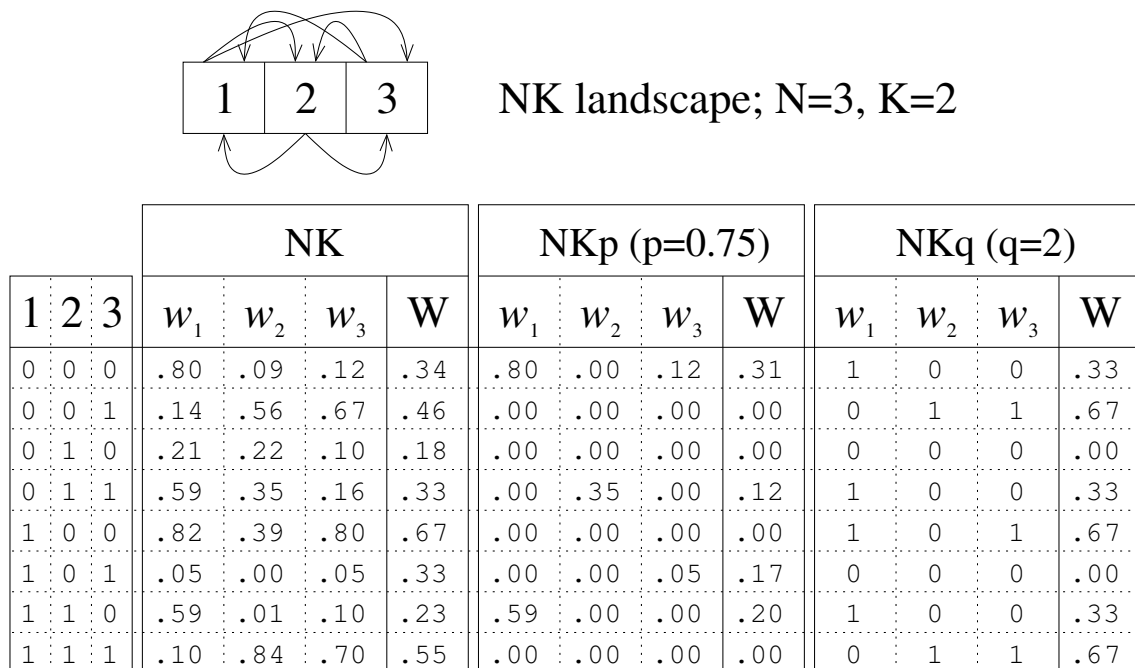


Figure 8.1: Example  $NK$ ,  $NKp$  and  $NKq$  landscapes. Top: organisms consist of three traits ( $N = 3$ ), each epistatically linked to all other traits ( $K = 2$ ). For each trait, there are  $2^{K+1}$  possible “input” combinations of traits. These are shown in the table, below. In the  $NK$  landscape, fitness contributions ( $w_i$ ) are assigned random real numbers: there is no neutrality; total fitness values ( $W$ ) are all different. The  $NKp$  landscape maps fitness contributions to 0.0 with probability  $p = 0.75$ : this produces neutrality. By quantising fitness contributions to integer values, the  $NKq$  landscape ( $q = 2$ ) produces three levels of neutrality.

$$W = \frac{1}{N(q-1)} \sum_i w_i \quad (8.1)$$

The fitness of all sequences thus falls in the range from zero to one, and there are  $Nq - N + 1$  possible fitness values in this range. In the limit, as  $q \rightarrow \infty$  the probability that two sequences will possess the same fitness becomes vanishingly small—equivalent to the  $NK$  model—and there is no neutrality. Neutrality is greatest when  $q$  takes the smallest possible value of two. Figure 8.1, right, demonstrates an example  $NKq$  landscape, with  $q = 2$ . This is a quantisation of the  $NK$  landscape, shown left. The  $NKq$  landscape has three distinct levels of neutrality, with fitness values 0.00, 0.33 and 0.67.

$NKp$  and  $NKq$  landscapes have inherent structural differences: while  $NKq$  landscapes correlate with the underlying  $NK$  landscape,  $NKp$  landscapes do not (Geard et al., 2002).  $NKq$  landscapes relate more closely to the intuitive idea of a “terraced”  $NK$  landscape in which rugged hillsides are flattened into local plateau. In contrast,  $NKp$  landscapes—with

neutrality largely confined to low regions of fitness space—more accurately capture the features of “lethal” mutations in biological organisms. Yet, despite structural differences, both  $NKp$  and  $NKq$  landscapes offer powerful, tuneable frameworks for investigating neutrality. However, since neither framework supports coevolutionary interactions, they are incapable of modelling coevolutionary neutrality.

To model coevolutionary systems using the  $NK$  framework, Kauffman and Johnsen (1991) proposed a generalization: the  $NKCS$  family of *coupled* fitness landscapes of different interacting species. Here, it is supposed that genes, or traits, in each species make fitness contributions which depend not only upon  $K$  other traits from within that species itself, but also upon  $C$  traits in each of the other species with which the species interacts. Thus, adaptations in one species may deform the landscape of its partners. Parameter  $S$  controls the total number of species that interact;  $S_i$  controls the (perhaps restricted) number of these with which any species,  $i$ , interacts.

In short, landscapes are coupled by expanding the random fitness table for each trait in species 2 such that it “looks” at its  $K$  internal epistatic inputs and also “looks” at the  $C$  external epistatic inputs from species 1 . . . This couples the two landscapes. Each is in the niche of the other (Kauffman & Johnsen, 1991).

The  $NKCS$  framework permits the study of how landscape ruggedness, landscape couplings, the number of species and the structure of the ecosystem affect coevolution. However, since the  $NKCS$  model randomly assigns the fitness contributions of each combination of trait inputs a real value—directly inherited from the original  $NK$  framework—there is no neutrality; hence,  $NKCS$  landscapes, despite allowing coevolutionary interactions, cannot model coevolutionary neutrality. In the following section, a novel framework for modelling coevolutionary neutrality is suggested.

## 8.2 Towards an $NKCqd$ Landscape

Considering the extensions to Kauffman’s  $NK$  framework introduced in the previous section, the necessary step needed to model coevolutionary neutrality seems clear: since  $NKp/q$  includes neutrality and  $NKCS$  includes coevolution, why not combine these frameworks? This can be simply achieved by quantising fitness contributions (using random integer values *a la*  $NKq$ ) in the  $NKCS$  framework; thus producing an  $NKCSq$  framework (or  $NKCq$ , given two populations) that incorporates both neutrality and coevolution.



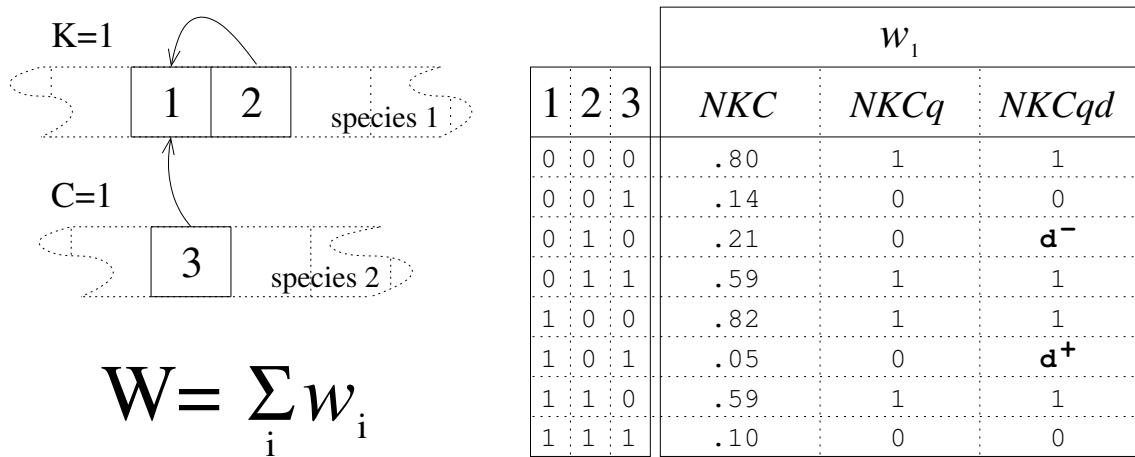


Figure 8.2: A partial  $NKCqd$  landscape ( $K = 1$ ,  $C = 1$ ,  $q = 2$ ,  $d = 0.25$  and  $N$  is unspecified). A fitness contribution of  $d$  indicates that traits have disengaged. If  $w_i = d^+$  then  $W = 1$ ; if  $w_i = d^-$  then  $W = 0$ . If no traits are disengaged,  $W$  is calculated as the mean value of all fitness contributions. Values of  $d^+$  and  $d^-$  are distributed equally.

As previously discussed, neutrality may be either inherent or transient. Inherent neutrality remains irrespective of environment: if two genotypes are inherently neutral, they either map to (near) identical phenotypes, or different phenotypes with (near) identical behaviour. Transient neutrality, however, varies over time as the environment changes: genotypes  $g_1$  and  $g_2$  may be (transiently) neutral at time  $t_0$ , but *not* at time  $t_1$ . The  $NKCq$  framework captures both.

While the  $NKCq$  framework models transient neutrality, it does not incorporate the widespread (catastrophic) neutrality that appears during periods of disengagement. Given the complexity of trait interactions (both within and between species)—represented by *random* fitness contributions—large scale transient neutrality is unlikely. To resolve this, an  $NKCqd$  (“disengaged”  $NKC$ ) generalization is proposed. Here, the  $NKC$  model is first quantised (in the manner of Newman & Engelhardt, 1998) to add neutrality, before a novel parameter,  $d$ , is used to control the probability of coevolutionary disengagement: if  $d = 0$  there is no disengagement; if  $d = 1.0$  the system is permanently disengaged.

To characterize disengagement, when particular coevolving traits disengage, the contributions of all other traits are ignored. In this way, large scale regions of transient neutrality form. As a biological analogy, suppose parasites find a strategy capable of killing all susceptible hosts. Then, as long as parasites and hosts maintain these traits, all other traits are irrelevant: size, metabolic rate, sexual attractiveness, etc.

Figure 8.2 displays a partial  $NKCqd$  landscape with unspecified  $N$ ,  $K = 1$  and  $C = 1$ . The table shows the fitness contributions of trait one,  $w_1$ , under each combination of input traits. Using Kauffman’s  $NKC$  framework, fitness values are real numbers: as a result,

there are no neutral mutations. Neutral mutations are introduced by quantising fitness contributions ( $q = 2$ ; middle column of table). In addition, disengagement is added using parameter  $d$ . For each input combination there is a probability,  $d$ , that disengagement will occur. In figure 8.2,  $d = 0.25$ ; however, for practical modelling  $d$  will usually take a much smaller value. When disengagement occurs in favour of a trait ( $d^+$ ), the entire organism (whole genotype) receives fitness  $W = 1$ , irrespective of the fitness contributions of all other traits. When disengagement occurs in favour of the opposition ( $d^-$ ), the entire organism receives fitness  $W = 0$ . Once a value of  $d$  has been read, the rest of the traits are ignored. If no traits are disengaged, the fitness of the entire organism,  $W$ , is calculated as the mean of all fitness contributions. In this way, the  $NKCqd$  landscape is able to model neutrality—both transient and inherent—and disengagement.

While this section has only suggested a potential framework, in future the  $NKCqd$  framework will be fully developed into a working model. Henceforth, it will be used to investigate the effects of neutrality and disengagement in coevolutionary systems.

### 8.3 Summary

In recent years, many researchers have become interested in the (beneficial) effects of neutrality. However, despite this enthusiasm, coevolutionary neutrality has been largely overlooked. While formal frameworks have been developed to study evolutionary neutrality, none have (as yet) been extended to incorporate coevolutionary neutrality. Here, steps have been taken towards the development of a framework with which to study coevolutionary neutrality and disengagement: the  $NKCqd$  landscape. In future, this framework will be fully developed into a working model. Hopefully, the insights it generates will enable greater understanding of coevolutionary phenomena; particularly disengaged drift.

# Chapter 9

## Conclusions

---

The following conclusions can be drawn from this thesis:

1. Disengagement is a common phenomenon that is more likely to occur when asymmetry favours one coevolving population. Though disengagement occurs when within population fitness diversity falls to zero, it is largely independent of genetic fitness diversity.
2. Disengagement can be combated by biasing selection in favour of individuals that win *some proportion*,  $\lambda$ , of competitions: the “reduced virulence” technique. This approach has been shown to outperform alternative techniques.
3. When utilising reduced virulence, there is a trade off between engagement and optimality: low virulence reduces the likelihood of disengagement, but sacrifices system performance. Thus, to maximise the quality of solutions, virulence should vary during the course of a run: when disengagement is unlikely,  $\lambda$  should be raised; when disengagement is likely,  $\lambda$  should be lowered.
4. CIAO plots are the most commonly used technique specifically designed for visualizing coevolutionary dynamics. There are few alternatives.
5. When used in isolation, CIAO plots suffer from potential ambiguity. To be interpreted with confidence, CIAO plots should be used as part of a suite of complementary techniques.

6. Steering tools enable users to interact with coevolutionary systems during run time. Such interactive visualization techniques offer the potential to generate novel insights into system behaviour, improve visualization and aid the development of new tools.

## 9.1 Summary of Contributions

This thesis has made the following contributions of new knowledge:

1. If an asymmetrical bias exists in favour of one coevolving population, “reducing the virulence” of the advantaged population—encouraging individuals to lose some competitions—reduces the likelihood of disengagement.
2. Potentially ambiguous, CIAO plots are difficult to interpret when used in isolation. As such, they should be used as part of a suite of complementary visualization techniques.
3. Interactive visualizations can be used to “steer” coevolutionary systems during run time.
4. The concepts of fitness propensity and fitness realization can be used to classify coevolution from non-coevolution in both artificial and biological systems.
5. Coevolutionary disengagement can be classified into three different categories: full disengagement, asymmetrical disengagement and collusive disengagement.
6. Linkage disequilibrium is a good measure of within population niching.
7. Image processing routines may aid the interpretation of CIAO plots.
8. CIAO plots can be used to visualize biological systems.

## 9.2 Limitations of the Thesis

The results presented in this thesis are subject to the following limitations:

**Reduced virulence:** Still in early stages of development, Dynamic virulence (chapter 6 is limited. The metric of engagedness ( $\epsilon$ ) and Dynamic virulence functions (Dynamic 1 and 2) are arbitrarily chosen.

**CIAO plots:** In chapter 7, CIAO plots are shown to be ambiguous with respect to irregular cycling and random drift. However, such classification involves subjective interpretation: plots may appear similar to one person and not another. There are currently no quantitative methods of classification.

**Coevolutionary steering:** The steering visualization introduced in chapter 7 is a demonstration model. Currently, no experimentation has been performed. While the potential of steering is hinted at, the ability of this system has not been quantifiably assessed.

### 9.3 Future Research

1. The possibility of a dynamically varying virulence function has been conceptually proven. However, this work is in its earliest stages of development. Firstly, a (less arbitrary) metric of engagement ( $\varepsilon$ ) needs to be formulated; perhaps utilising an *entropy* measure of the results of pairwise contests. Secondly, the relationship between engagement and virulence (the dynamic function) needs much work. Specifically, should dynamic virulence be sensitive to particular values of  $\varepsilon$ ? Should virulence be calculated from the current value of  $\varepsilon$ , or should previous values be considered, providing evolutionary momentum? These lines of research are likely to be time consuming, but offer potentially high rewards.
2. What happens at the moment of re-engagement? It has been shown that high virulence increases the likelihood of populations “bouncing apart” once they re-engage. However, it would be useful to have a better understanding of this. Under different virulence schemes, it would be useful to visualize each individual lineage for several generations immediately following re-engagement. As systems immediately following re-engagement are the most sensitive to virulence—in the sense that disengagement is most likely to occur—analysing these periods offer the best chance of gaining insight into the relationship between virulence, thrashing and disengagement. Such detailed micro-analysis of the effects of virulence may generate insight useful for developing dynamic virulence routines.
3. The coevolutionary steering platform developed in chapter 7 can be used to explore disengagement and its relationship with virulence. To tease out tacit knowledge of users, novices and experts will attempt to maximise the absolute fitness  $n$ —number of 1s—of both populations by varying virulence during runtime as asymmetry varies. If experts perform best, it is likely that they are using some (tacit)

knowledge unavailable to novices. Once this knowledge has been accessed, there is potential to incorporate it procedurally into a autonomously varying dynamic virulence function. Alternatively, if such knowledge *cannot* be operationally defined, the experiments will serve to demonstrate the utility of steering; not only as a research tool, but as a technique capable of improving the performance of coevolution.

4. Chapter 8 suggested a potential framework—the *NKCqd* model—to study coevolutionary neutrality. Further development of this (or a similar) framework will greatly benefit coevolutionary computing.
5. This thesis has suggested that funnelling (refer to section 7.5), a previously unobserved phenomenon, has potential to offer a “scaffold” for coevolutionary development. This could be proven by using a simple domain in which coevolution cannot reach the global optimum without funnelling. If proven, the scaffolding ability of funnelling is potentially very powerful.
6. Fourier analysis of CIAO plots offers a potential framework for quantifying cycling within a coevolutionary system. If successful, this approach may allow the development of a cycling “metric”. When used in connection with a metric of engagedness, these parameters may allow users to observe the “healthiness” of a coevolutionary system—whether the system is exhibiting “problematic” dynamics, or not—during run time (assuming CIAO plots are updated each generation, rather than calculated once a run has terminated). Further, by incorporating steering, when such problematic dynamics occur, users may take preventative action.

## 9.4 Verification

Some of the results presented in this thesis have been independently verified. Reduced virulence has been shown to improve performance when coevolving simulated maze-solving agents and mazes-to-be-solved: without a reduction in virulence, mazes quickly become unsolvable (Clarke, 2002). The sorting network results presented in chapter 6 have also been reproduced.<sup>1</sup>

---

<sup>1</sup>Simon McGregor, School of Cognitive and Computing Sciences, University of Sussex, UK, 2003 (personal communication).

# Bibliography

- Adami, C., Belew, R., Kitano, H., & Taylor, C. (Eds.). (1998). *Artificial Life VI*. MIT Press, Cambridge MA.
- Angeline, P. J. (1994). An alternate interpretation of the iterated prisoner's dilemma and the evolution of non-mutual cooperation. In Brooks & Maes (1994), pp. 353–358.
- Angeline, P. J., & Pollack, J. B. (1993). Competitive environments evolve better solutions for complex tasks. In Forrest (1993), pp. 264–270.
- Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In Davis, L. (Ed.), *Genetic Algorithms and Simulated Annealing*, pp. 32–41. Morgan Kaufmann, Los Altos CA.
- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York NY.
- Bäck, T. (1996). *Evolutionary Algorithms in Theory and Practice: Evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, Oxford, UK.
- Baker, J. E. (1985). Adaptive selection methods for genetic algorithms. In Grefenstette (1985), pp. 101–111.
- Baker, M. (2002a). 'Murder' of the national curriculum. *BBC News Online*, 9th February, UK.
- Baker, M. (2002b). What makes teachers cheat? *BBC News Online*, 15th June, UK.
- Bakker, R. T. (1983). The deer flees, the wolf pursues: Incongruencies in predator-prey coevolution. In Futuyma & Slatkin (1983), pp. 350–382.
- Barnett, L. (1998). Ruggedness and neutrality—the NKp family of fitness landscapes. In Adami et al. (1998), pp. 18–27.

- Barrett, J. A. (1983). Plant-fungus symbioses. In Futuyma & Slatkin (1983), pp. 137–160.
- Barry, A. M. (Ed.). (2002). *Bird of a Feather Workshops, Genetic and Evolutionary Computation Conference*. AAAI, New York NY.
- Barton, N. H., & Gale, K. S. (1993). Genetic analysis of hybrid zones. In Harrison, R. G. (Ed.), *Hybrid Zones and the Evolutionary Process*, pp. 13–45. Oxford University Press, Oxford, UK.
- Bates, H. W. (1862). Contributions to an insect fauna of the amazon valley. *Transactions of the Linnean Society of London*, 23, 495–566.
- Bedau, M. A., & Packard, N. H. (1992). Measurement of evolutionary activity, teleology, and life. In Langton et al. (1992), pp. 431–461.
- Bedau, M. (1996). The nature of life. In Boden, M. A. (Ed.), *The Philosophy of Artificial Life*, pp. 332–357. Oxford University Press, Oxford, UK.
- Bedau, M. A., & Brown, C. (1998). Visualizing evolutionary activity of genotypes. Technical report 98-03-23, Santa Fe Institute, Santa Fe NM.
- Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1), 91–122.
- Bosman, P. A. N. (1999). EA visualizer tutorial. Technical report UU-CS-1999-20, Utrecht University.
- Brooks, R. A., & Maes, P. (Eds.). (1994). *Artificial Life IV*. MIT Press, Cambridge MA.
- Brown, jr., W. L., & Wilson, E. O. (1956). Character displacement. *Systematics Zoology*, 7, 49–64.
- Bucci, A., & Pollack, J. B. (2003). Focusing versus intransitivity: Geometrical aspects of co-evolution. In Cantú-Paz, E., Foster, J. A., Deb, K., Davis, D., Roy, R., O'Reilly, U.-M., Beyer, H.-G., Standish, R., Kendall, G., Wilson, S., Harman, M., Wegener, J., Dasgupta, D., Potter, M. A., Schultz, A. C., Dowsland, K., Jonoska, N., & Miller, J. (Eds.), *Genetic and Evolutionary Computation Conference*, pp. 250–261. Springer-Verlag, Berlin, Germany.
- Bucci, A., & Pollack, J. (2002). Order-theoretic analysis of coevolution problems: Co-evolutionary statics. In Barry (2002), pp. 229–235.



- Bullock, S. (1995). Co-evolutionary design: Implications for evolutionary robotics. Cognitive science research paper 384, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.
- Bullock, S. (1999). Are artificial mutation biases unnatural? In Floreano, D., Nicoud, J.-D., & Mondada, F. (Eds.), *Fifth European Conference on Artificial Life*, pp. 64–73. Springer-Verlag, Berlin, Germany.
- Bullock, S. (2001). Smooth operator? Understanding and visualising mutation bias. In Kelemen & Sosik (2001), pp. 602–612.
- Bullock, S. (2002). Will selection for mutational robustness significantly retard evolutionary innovation on neutral networks? In Standish, R., Bedau, M., & Abbass, H. (Eds.), *Artificial Life VIII*, pp. 192–201. MIT Press, Cambridge MA.
- Bullock, S., Cartlidge, J., & Thompson, M. (2002). Prospects for computational steering of evolutionary computation. In Bilotta, E., Groß, D., Smith, T., Lenaerts, T., Bullock, S., Lund, H. H., Bird, J., Watson, R., Pantano, P., Pagliarini, L., Abbass, H., Standish, R., & Bedau, M. (Eds.), *Workshop Proceedings, Artificial Life VIII*, pp. 131–137. MIT Press, Cambridge MA.
- Cartlidge, J., & Bullock, S. (2002). Learning lessons from the common cold: How reducing parasite virulence improves coevolutionary optimization. In *Congress on Evolutionary Computation*, pp. 1420–1425. IEEE Press, Piscataway NJ.
- Cartlidge, J., & Bullock, S. (2003). Caring versus sharing: How to maintain engagement and diversity in coevolving populations. In Banzhaf, W., Christaller, T., Dittrich, P., Kim, J. T., & Ziegler, J. (Eds.), *Seventh European Conference on Artificial Life*, pp. 299–308. Springer-Verlag, Berlin, Germany.
- Cartlidge, J., & Bullock, S. (2004a). Combating coevolutionary disengagement by reducing parasite virulence. *Evolutionary Computation*, 12(2), 193–222.
- Cartlidge, J., & Bullock, S. (2004b). Unpicking tartan CIAO plots: Understanding irregular coevolutionary cycling. *Adaptive Behaviour*, 12(2), 69–92.
- Chalmers, A. F. (1999). *What Is This Thing Called Science?* (Third edition). Open University Press, Milton Keynes, UK.
- Clarke, P. (2002). *Evolving Maze-solving Agents*. Unpublished B.Sc. dissertation, School of Computing, University of Leeds, UK.

- Cliff, D., & Miller, G. F. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Morán et al. (1995), pp. 200–218.
- Cliff, D., & Miller, G. F. (1996). Coevolution of pursuit and evasion II: Simulation methods and results. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J. B., & Wilson, S. W. (Eds.), *Fourth International Conference on Simulation of Adaptive Behavior*, pp. 506–515. MIT Press, Cambridge MA.
- Collins, J. J. (1999). Visualization of evolutionary algorithms using principal component analysis. In Wu, A. S. (Ed.), *Evolutionary Computation Visualization Workshop, Genetic and Evolutionary Computation Conference*. Morgan Kaufmann, San Francisco CA.
- Collins, T. D. (1997). Using software visualization technology to help evolutionary algorithm users validate their solutions. In Bäck, T. (Ed.), *Seventh International Conference on Genetic Algorithms*, pp. 307–314. Morgan Kaufmann, San Francisco CA.
- Collins, T. D. (1998). Understanding evolutionary computing: A hands on approach. In *IEEE International Conference on Evolutionary Computation*, pp. 564–569. IEEE Press, Piscataway NJ.
- Combes, C. (1991). Ethological aspects of parasite transmission. *The American Naturalist*, 138(4).
- Dabs, T., & Schoof, J. (1995). A graphical user interface for genetic algorithms. Technical report 98, Lehrstuhl für Informatik II, University of Würzburg, Germany.
- Darwen, P., & Yao, X. (1995). On evolving robust strategies for iterated prisoner's dilemma. In Yao, X. (Ed.), *Progress in Evolutionary Computation*, pp. 276–292. Springer-Verlag, Heidelberg, Germany.
- Darwen, P., & Yao, X. (1996). Every niching method has its niche: Fitness sharing and implicit sharing compared. In Voigt, H.-M., Ebeling, W., Rechenberg, I., & Schwefel, H.-P. (Eds.), *Parallel Problem Solving from Nature IV*, pp. 398–407. Springer-Verlag, Berlin, Germany.
- Darwen, P., & Yao, X. (2001). Why more choices cause less cooperation in iterated prisoner's dilemma. In *Congress on Evolutionary Computation*, pp. 987–994. IEEE Press, Piscataway NJ.

- Darwen, P. J., & Yao, X. (1997). Speciation as automatic categorical modularization. *IEEE Transactions on Evolutionary Computation*, 1(2), 101–108.
- Darwen, P. J., & Yao, X. (2000). Does extra genetic diversity maintain escalation in a co-evolutionary arms race. *International Journal of Knowledge-Based Intelligent Engineering Systems*, 4(3), 191–200.
- Darwin, C. (1859). *On the Origin of Species by Means of Natural Selection or the Preservation of Favoured Races in the Struggle for Life*. John Murray, London.
- Darwin, C. (1862). *The Various Contrivances by Which Orchids Are Fertilized by Insects*. John Murray, London.
- Darwin, C. (1985). *On the Origin of Species by Means of Natural Selection or the Preservation of Favoured Races in the Struggle for Life*. Edited with an introduction by J. W. Burrow. Penguin Classics, London. First published as Darwin (1859).
- Davis, L. (Ed.). (1991). *Handbook of genetic algorithms*. Van Nostrand Reinhold, New York NY.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, Oxford, UK.
- Dawkins, R. (1982). *The Extended Phenotype*. Oxford University Press, Oxford, UK.
- Dawkins, R. (1986). *The Blind Watchmaker*. Longman, Harlow, UK.
- Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London, Series B*, 205, 489–511.
- De Jong, E. D., & Pollack, J. B. (2004). Ideal evaluation from coevolution. *Evolutionary Computation*, 11(2).
- De Jong, K. A. (1975). *An analysis of the behavior of a class of genetic adaptive systems*. Unpublished Ph.D. thesis, University of Michigan, Michigan MI.
- De Jong, K. A. (1985). Genetic algorithms: A ten year perspective. In Grefenstette (1985), pp. 169–177.
- Dorin, A. (2001). Aesthetic fitness and artificial evolution for the selection of imagery from the mythical infinite library. In Kelemen & Sosík (2001), pp. 659–668.

- Downes, S., & Shine, R. (1998). Sedentary snakes and gullible geckos: Predator-prey coevolution in nocturnal rock-dwelling reptiles. *Animal Behaviour*, 55(5), 1373–1385.
- Dybowski, R., Collins, T. D., & Weller, P. (1996). Visualization of binary string convergence by sammon mapping. In Fogel, L., Angeline, P., & Bäck, T. (Eds.), *Fifth Annual Conference on Evolutionary Programming*, pp. 377–383. MIT Press, Cambridge MA.
- Dzwinel, W. (1994). How to make Sammon's mapping useful for multidimensional data structure analysis. *Pattern Recognition*, 27, 949–959.
- Ehrlich, P. R., & Raven, P. H. (1964). Butterflies and plants: A study in co-evolution. *Evolution*, 18, 586–608.
- Ehrman, L. (1983). Endosymbiosis. In Futuyma & Slatkin (1983), pp. 128–136.
- Eldredge, N., & Gould, S. J. (1973). Punctuated equilibria: An alternative to phyletic gradualism. In Schopf, T. J. M. (Ed.), *Models in Paleobiology*, pp. 82–115. Freeman, Cooper and Co, San Francisco CA.
- Eshelman, L. J. (Ed.). (1995). *Sixth International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo CA.
- Falconer, D. S. (1996). *Introduction to Quantitative Genetics*. Longman Group, London, UK.
- Fenner, F., & Ratcliffe, F. N. (1965). *Myxomatosis*. Cambridge University Press, Cambridge, UK.
- Ficici, S. G., & Pollack, J. B. (1998a). Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In Adami et al. (1998), pp. 238–247.
- Ficici, S. G., & Pollack, J. B. (1998b). Coevolving communicative behavior in a linear pursuer-evader game. In Pfeifer, R., Blumberg, B., Meyer, J.-A., & Wilson, S. W. (Eds.), *Fifth International Conference on Simulation of Adaptive Behavior*, pp. 505–514. MIT Press, Cambridge MA.
- Ficici, S. G., & Pollack, J. B. (2001). Pareto optimality in coevolutionary learning. In Kelemen (2001), pp. 316–325.

- Ficici, S. G., & Pollack, J. B. (2003). A game-theoretic memory mechanism for co-evolution. In Cantú-Paz, E., Foster, J. A., Deb, K., David, L., Roy, R., O'Reilly, U.-M., Beyer, H.-G., Standish, R., Kendall, G., Wilson, S., Harman, M., Wegener, J., Dasgupta, D., Potter, M. A., Schultz, A. C., Dowsland, K. A., Jonoska, N., & Miller, J. (Eds.), *Genetic and Evolutionary Computation Conference*, pp. 286–297. Springer-Verlag, Berlin, Germany.
- Flor, H. H. (1955). Host-parasite interaction in flax rust—its genetics and other implications. *Phytopathology*, *45*, 680–685.
- Floreano, D., & Nolfi, S. (1997a). Adaptive behaviour in competing co-evolving species. In Hisbands, P., & Harvey, I. (Eds.), *Fourth European Conference on Artificial Life*, pp. 378–387. MIT Press, Cambridge MA.
- Floreano, D., & Nolfi, S. (1997b). God save the red queen! Competition in co-evolutionary robotics. In Koza, J. R., Deb, K., Dorigo, M., Fogel, D. B., Garzon, M., Iba, H., & Riolo, R. L. (Eds.), *Second Annual Conference on Genetic Programming*, pp. 398–406. Morgan Kaufmann, San Francisco CA.
- Floreano, D., Nolfi, S., & Mondada, F. (1998). Competitive co-evolutionary robotics: From theory to practice. In Pfeifer (1998), pp. 512–524.
- Fogel, D. B. (2000). Evolving a checkers player without relying on human expertise. *Intelligence*, *11*(2), 20–27.
- Fogel, L. J., Owens, A. J., & Walsh, M. J. (1966). *Artificial Intelligence through Simulated Evolution*. John Wiley, New York NY.
- Fontana, W., & Schuster, P. (1998a). Continuity in evolution: On the nature of transitions. *Science*, *280*, 1451–1455.
- Fontana, W., & Schuster, P. (1998b). Shaping space: The possible and the attainable in rna genotype-phenotype mapping. *Journal of Theoretical Biology*, *194*, 491–515.
- Forrest, S. (Ed.). (1993). *Fifth International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo CA.
- Funes, P., Sklar, E., Juillé, H., & Pollack, J. B. (1998). Animal - animat coevolution: Using the animal population as fitness function. In Rolf Pfeifer & Wilson (1998), pp. 525–533.

- Funes, P. (2001). *Evolution of Complexity in Real-World Domains*. Unpublished Ph.D. thesis, Brandeis University, Waltham MA.
- Futuyma, D. J. (1979). *Evolutionary Biology*. Sinauer Associates, Sunderland MA.
- Futuyma, D. J., & Slatkin, M. (Eds.). (1983). *Coevolution*. Sinauer Associates, Sunderland MA.
- Gale, J. S., & Eaves, L. J. (1975). Logic of animal conflict. *Nature*, 254, 463–464.
- Gallun, R. L. (1977). The genetic basis of hessian fly epidemics. *Annual New York Academy Science*, 287, 223–229.
- Geard, N., Wiles, J., Hallinan, J., Tonkes, B., & Skellett, B. (2002). Comparison of neutral landscapes—NK, NKp and NKq. In *Congress on Evolutionary Computation*, pp. 205–210. IEEE Press, Piscataway NJ.
- Gilbert, L. E., & Ehrlich, P. R. (Eds.). (1975). *Coevolution of Animals and Plants*. University of Texas Press, Austin TX.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading MA.
- Goldberg, D. E., & Richardson, J. J. (1987). Genetic algorithms with sharing for multimodal function optimization. In Grefenstette, J. J. (Ed.), *Second International Conference on Genetic Algorithms*, pp. 41–49. Lawrence Erlbaum Associates, Hillsdale NJ.
- Gould, S. J., & Eldredge, N. (1977). Punctuated equilibria: The tempo and mode of evolution reconsidered. *paleobiology*, 3, 115–151.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marcos and the Panglossian program: A critique of the adaptationist programme. *Proceedings of the Royal Society of London, Series B*, 205, 281–288.
- Grafen, A. (1991). Modelling in behavioural ecology. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach*, pp. 5–31. Blackwell Scientific Publications, Oxford, UK.
- Grefenstette, J. J. (Ed.). (1985). *First International Conference on Genetic Algorithms*. Lawrence Erlbaum Associates, Hillsdale NJ.

- Gruau, F. (1995). Automatic definition of modular neural networks. *Adaptive Behaviour*, 3(2), 151–183.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour, I and II. *Journal of Theoretical Biology*, 7, 1–16; 17–32.
- Hamilton, W. D. (1980). Sex versus non-sex versus parasite. *OIKOS*, 35, 282–290.
- Harrison, A. (2003). Are A-levels getting easier? *BBC News Online*, 14th August, UK.
- Hart, E., & Ross, P. (2001). GAVEL: A new tool for genetic algorithm visualization. *IEEE Transactions on Evolutionary Computation*, 5(4), 335–348.
- Harvey, I. (1993). The puzzle of the persistent question marks: A case study in genetic drift. In Forrest (1993), pp. 15–22.
- Harvey, I., & Thompson, A. (1996). Through the labyrinth evolution finds a way: A silicon ridge. In Higuchi, T., Iwata, M., & Wexlin, L. (Eds.), *First International Conference on Evolvable Systems*, pp. 406–422. Springer-Verlag, Berlin, Germany.
- Hauert, C., De Monte, S., Hofbauer, J., & Sigmund, K. (2002). Volunteering as red queen mechanism for cooperation in public goods games. *Science*, 296, 1129–1132.
- Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42, 228–234.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495–502.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor MI.
- Holland, J. H. (1992). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. MIT Press/Bradford Books edition. Reprint of Holland (1975).
- Hood, G. M. (1997). Spatial structure in a host-pathogen system. In Klomp, N. I., & Lunt, I. (Eds.), *Frontiers in Ecology: Building the Links*, pp. 265–275. Elsevier Science, Oxford, UK.
- Hornby, G. S., Lipson, H., & Pollack, J. B. (2001). Evolution of generative design systems for modular physical robots. In *International Conference on Robotics and Automation*. IEEE Press, Piscataway, NJ.

- Hornby, G. S., & Pollack, J. B. (2001). Body-brain coevolution using l-systems as a generative encoding. In Spencer et al. (2001), pp. 868–875.
- Hornby, G. S., & Pollack, J. B. (2002). Creating high-level components with a generative representation for body-brain evolution.. *Artificial Life*, 8(3).
- Husbands, P., & Mill, F. (1991). Simulated co-evolution as the mechanism for emergent planning and scheduling. In Belew, R. K., & Booker, L. B. (Eds.), *Fourth International Conference on Genetic Algorithms*, pp. 264–270. Morgan Kaufman, San Mateo CA.
- Huxley, J. (1942). *Evolution, the modern synthesis*. Allen and Unwin, London.
- Huynen, M. (1996). Exploring phenotype space through neutral evolution. *Journal of Molecular Evolution*, 43, 165–169.
- Huynen, M., Stadler, P., & Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Science USA*, 93(1), 397–401.
- Janzen, D. H. (1980). When is it co-evolution? *Evolution*, 34(3), 611–612.
- Janzen, D. H. (1966). Coevolution of mutualism between ants and acacias in central america. *Evolution*, 20, 249–275.
- Juillé, H. (1995). Incremental co-evolution of organisms: A new approach for optimisation and discovery of strategies. In Morán et al. (1995), pp. 246–260.
- Juillé, H., & Pollack, J. B. (1998a). Coevolutionary learning: A case study. In Shavlik, J. W. (Ed.), *Fifteenth International Conference on Machine Learning*. Morgan Kaufmann, San Francisco CA.
- Juillé, H., & Pollack, J. B. (1998b). Coevolving the “ideal” trainer: Application to the discovery of cellular automata rules. In Koza, J. R., Banzhaf, W., Chellapilla, K., Deb, K., Dorigo, M., Fogel, D. B., Garzon, M. H., Goldberg, D. E., Iba, H., & Riolo, R. (Eds.), *Third Annual Conference on Genetic Programming*, pp. 519–527. Morgan Kaufmann, San Francisco CA.
- Juille, H., & Pollack, J. B. (1996). Co-evolving intertwined spirals. In Fogel, L. J., Angeline, P. J., & Baeck, T. (Eds.), *Fifth Annual Conference on Evolutionary Programming*, pp. 461–468. MIT Press, Cambridge, MA.



- Kapsalis, A., & Smith, G. (1992). *The GAmeter Development Toolkit User Interface Manual*. University of East Anglia, Norwich, UK.
- Kauffman, S. A. (1993). *The Origins of Order: Organization and Selection in Evolution*. Oxford University Press, Oxford, UK.
- Kauffman, S. A., & Johnsen, S. (1991). Co-evolution at the edge of chaos: Coupled fitness landscapes, poised states, and co-evolutionary avalanches. *Journal of Theoretical Biology*, 149(4), 467–505.
- Kelemen, J., & Sosík, P. (Eds.). (2001). *Sixth European Conference on Artificial Life*. Springer-Verlag, Berlin, Germany.
- Kerr, B., Riley, M. A., Feldman, M. W., & Bohannan, B. J. M. (2002). Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature*, 418(11), 171–174.
- Knuth, D. E. (1973). *The Art of Computer Programming - Sorting and Searching*, Vol. 3. Addison Wesley, Reading MA.
- Knuth, D. E. (1998). *The Art of Computer Programming - Sorting and Searching* (2 edition), Vol. 3. Addison Wesley Longman, Redwood City CA.
- Koza, J. R. (Ed.). (1993). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge MA.
- Koza, J. R. (1991). Evolution and co-evolution of computer programs to control independent-acting agents. In Meyer, J.-A., & Wilson, S. W. (Eds.), *First International Conference on Simulation of Adaptive Behavior*, pp. 366–375. MIT Press, Cambridge, MA.
- Kuhn, T. (1970). *The Structure of Scientific Revolutions*. University of Chicago, Chicago IL.
- Lack, D. (1947). *Darwin's Finches*. Cambridge University Press, Cambridge, UK.
- Lamarck, J.-B. (1809). *Philosophie zoologique, ou exposition des considérations relatives à l'histoire naturelle des animaux*. E. Dentu, Paris, France.
- Land, M., & Belew, R. K. (1995). No perfect two-state cellular automata for density classification exists. *Physical Review Letters*, 74(25), 1548–1550.

- Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.). (1992). *Artificial Life II*. Addison-Wesley, Redwood City CA.
- Lawlor, L. R., & Maynard Smith, J. (1976). The coevolution and stability of competing species. *American Naturalist*, *110*, 79–99.
- Levin, B. R. (1996). The evolution and maintenance of virulence in microparasites. *Emerging Infectious Diseases*, *2*(2).
- Lewontin, R. C. (1970). The units of selection. *Annual Review of Ecology and Systematics*, *1*, 1–14.
- Li, S., & Loew, M. H. (1987). The quadcode and its arithmetic. *Communications of the ACM*, *30*(7), 621–626.
- Lindenmayer, A. (1968). Mathematical models for cellular interaction in development: Parts I and II. *Journal of Theoretical Biology*, *18*, 280–299 and 300–315.
- Lindgren, K., & Nordahl, M. G. (1994). Evolutionary dynamics of spatial games. *Physica D*, *75*, 292–309.
- Lubberts, A., & Miikkulainen, R. (2001). Co-evolving a go-playing neural network. In Belew, R. K., & Juillè, H. (Eds.), *Workshop on Learning, Adaptation and Approximation in Evolutionary Computation, Genetic and Evolutionary Computation Conference*, pp. 14–19.
- Luke, S., & Wiegand, R. P. (2002). When coevolutionary algorithms exhibit evolutionary dynamics. In Barry (2002), pp. 236–241.
- Machiavelli, N. (1998). *The Prince*. Edited with an Introduction and Notes by P. Bondanella; Translated by P. Bondanella and M. Musa. Oxford University Press, Oxford, UK.
- Mahfoud, S. W. (1995). *Niching Methods for Genetic Algorithms*. Unpublished Ph.D. thesis, University of Illinois at Urbana-Champaign, Urbana IL.
- Maley, C. C. (1994). A model of the effects of dispersal distance on the evolution of virulence in parasites. In Brooks & Maes (1994), pp. 102–108.
- Malthus, T. R. (1798). *An Essay on the Principle of Population, as It Affects the Future Improvement of Society*. J. Johnson, London.

- Margulis, L. (1970). *Origin of Eukaryotic Cells*. Yale University Press, New Haven CT.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W H Freeman, San Francisco CA.
- May, R. M., & Anderson, R. M. (1983). Parasite-host coevolution. In Futuyma & Slatkin (1983), pp. 186–206.
- Maynard Smith, J. (1989). *Evolutionary Genetics*. Oxford University Press, Oxford, UK.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- Maynard Smith, J. (1996). The games lizards play. *Nature*, 380(6571), 198–199.
- Maynard Smith, J., & Price, G. R. (1973). The logic of animal conflict. *Nature*, 246, 15–18.
- Mayr, E. (1982). *The Growth of Biological Thought: Diversity, Evolution and Inheritance*. Harvard University Press, Cambridge MA.
- Mayr, E. (2001). *What Evolution Is*. Basic Books, New York NY.
- Mayr, E., & Provine, W. B. (Eds.). (1980). *The Evolutionary Synthesis: Perspectives on the Unification of Biology*. Harvard University Press, Cambridge MA.
- Miller, G. F., & Cliff, D. (1994a). Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations. Technical report CSRP311, University of Sussex, Brighton, UK, School of Cognitive and Computing Sciences.
- Miller, G. F., & Cliff, D. (1994b). Protean behaviour in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *Third International Conference on Simulation of Adaptive Behavior*, pp. 411–420. MIT Press, Cambridge MA.
- Mills, S. K., & Beatty, J. H. (1994). The propensity interpretation of fitness. In Sober, E. (Ed.), *Conceptual Issues in Evolutionary Biology* (Second edition), pp. 3–24. MIT Press, Cambridge MA.
- Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge MA.
- Mitchell, M., Crutchfield, J. P., & Hraber, P. T. (1994). Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D*, 75, 361–391.

- Mitter, C., & Brooks, D. R. (1983). Phylogenetic aspects of coevolution. In Futuyma & Slatkin (1983), pp. 65–98.
- Mode, C. J. (1958). A mathematical model for the co-evolution of obligate parasites and their hosts. *Evolution*, 12, 158–165.
- Moore, G. E. (1903). *Principia Ethica*. Cambridge University Press, Cambridge, UK.
- Morán, F., Moreno, A., Merelo, J. J., & Chacón, P. (Eds.). (1995). *Third European Conference on Artificial Life*. Springer-Verlag, Berlin, Germany.
- Müller, F. (1879). Ituna and thyridia: A remarkable case of mimicry in butterflies..
- Newman, M. E. J., & Engelhardt, R. (1998). Effect of neutral selection on the evolution of molecular species. *Proceedings of the Royal Society of London, Series B*, 256, 1333–1338.
- Noble, J., & Watson, R. A. (2001). Pareto coevolution: Using performance against co-evolved opponents in a game as dimensions for pareto selection. In Spencer et al. (2001), pp. 493–500.
- Nolfi, S., & Floreano, D. (1998). How co-evolution can enhance the adaptive power of artificial evolution: Implications for evolutionary robotics. In Husbands, P., & Meyer, J.-A. (Eds.), *First European Workshop on Evolutionary Robotics*, pp. 22–38. Springer-Verlag, Berlin, Germany.
- Nordenskiöld, E. (1935). *The History of Biology*. Alfred A. Knopf, New York NY.
- Olsson, B. (1996). Optimisation using a host-parasite model with variable-size distributed populations. In *International Conference on Evolutionary Computation*, pp. 295–299. IEEE Press, Piscataway NJ.
- Olsson, B. (2001). Co-evolutionary search in asymmetric spaces. *Information Sciences*, 133(3), 103–125.
- Østergaard, E. H., & Lund, H. (2003). Co-evolving complex robot behavior. In Tyrrell, A. M., Haddow, P. C., & Torresen, J. (Eds.), *Fifth International Conference on Evolvable Systems*, Vol. 2606, pp. 308–319. Springer-Verlag, Berlin, Germany.
- Pagie, L., & Mitchell, M. (2002). A comparison of evolutionary and coevolutionary search. *International Journal of Computational Intelligence and Applications*, 2(1), 53–69.

- Pagie, L., & Hogeweg, P. (1998). Evolutionary consequences of coevolving targets. *Evolutionary Computation*, 5(4), 401–418.
- Pagie, L., & Hogeweg, P. (2000). Information integration and red queen dynamics in coevolutionary optimization. In *Congress on Evolutionary Computation*, pp. 1260–1267. IEEE Press, Piscataway NJ.
- Paredis, J. (1995). The symbiotic evolution of solutions and their representations. In Eshelman (1995), pp. 359–365.
- Paredis, J. (1997). Coevolving cellular automata: Be aware of the red queen! In Bäch, T. (Ed.), *Seventh International Conference on Genetic Algorithms*, pp. 393–399. Morgan Kaufmann, San Francisco CA.
- Parker, G. (1979). Sexual selection and sexual conflict. In Blum, M. S., & Blum, N. B. (Eds.), *Sexual Selection and Reproductive Competition in Insects*, pp. 123–166. Academic Press, New York NY.
- Parker, G., & Macnair, M. (1979). Models of parent-offspring conflict IV. Suppression: Evolutionary retaliation by the parent. *Animal Behaviour*, 27, 1210–1235.
- Parker, S. G., Johnson, C. R., & Beazley, D. (1997). Computational steering software systems and strategies. *IEEE Computational Science & Engineering*, 4(4), 50–59.
- Pfeifer, R., Blumberg, B., Meyer, J.-A., & Wilson, S. W. (Eds.). (1998). *Fifth International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge MA.
- Pohlheim, H. (1994-2001). GEATbx: Genetic and evolutionary algorithm toolbox for matlab. <http://www.geatbx.com/>.
- Pohlheim, H. (1999). Visualization of evolutionary algorithms: Set of standard techniques and multidimensional visualization. In Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, V., Jakiela, M., & Smith, R. E. (Eds.), *Genetic and Evolutionary Computation Conference*, pp. 533–540. Morgan Kaufmann, San Francisco CA.
- Pollack, J. B., Blair, A., & Land, M. (1996). Coevolution of a backgammon player. In Langton, C. G., & Shimohara, T. (Eds.), *Artificial Life V*, pp. 92–98. MIT Press, Cambridge MA.

- Pollack, J. B., & Blair, A. D. (1998). Co-evolution in the successful learning of a backgammon strategy. *Machine Learning*, 32(3), 225–240.
- Potter, M. A., & De Jong, K. A. (1994). A cooperative coevolutionary approach to function optimisation. In Davidor, Y., Schwefel, H.-P., & Männer, R. (Eds.), *Parallel Problem Solving from Nature III*, pp. 249–257. Springer-Verlag, Berlin, Germany.
- Potter, M. A., & De Jong, K. A. (1995). Evolving neural networks with collaborative species. In Oren, T. I., & Birta, L. G. (Eds.), *Summer Computer Simulation Conference*, pp. 340–345. The Society of Computer Simulation, San Diego CA.
- Potter, M. A., & De Jong, K. A. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8(1), 1–29.
- Puppala, N., Sen, S., & Gordin, M. (1998). Shared memory based cooperative coevolution. In Fogel, D. B. (Ed.), *International Conference on Evolutionary Computation*, pp. 570–574. IEEE Press, Piscataway NJ.
- Rechenberg, I. (Ed.). (1973). *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Frommann-Holzboorg, Stuttgart.
- Reed, J., Toombs, R., & Barricelli, N. A. (1967). Simulation of biological evolution and machine learning: I. Selection of self-reproducing numeric patterns by data processing machines, effects of hereditary control, mutation type and crossing. *Journal of Theoretical Biology*, 17, 319–342.
- Reynolds, C. W. (1994). Competition, coevolution and the game of tag. In Brooks & Maes (1994), pp. 59–69.
- Rosin, C. D. (1997). *Coevolutionary Search among Adversaries*. Unpublished Ph.D. thesis, Department of Computer Science, University of California, San Diego CA.
- Rosin, C. D., & Belew, R. K. (1995). Methods for competitive co-evolution: Finding opponents worth beating. In Eshelman (1995), pp. 373–380.
- Rosin, C. D., & Belew, R. K. (1997). New methods for competitive coevolution. *Evolutionary Computation*, 5(1), 1–29.
- Roughgarden, J. (1983a). Coevolution between competitors. In Futuyma & Slatkin (1983), pp. 383–403.

- Roughgarden, J. (1983b). The theory of coevolution. In Futuyma & Slatkin (1983), pp. 33–64.
- Sammon, J. (1969). A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers*, C-18(5), 401–408.
- Semmann, D., Krambeck, H.-J., & Millinski, M. (2003). Volunteering leads to rock-paper-scissors dynamics in a public goods game. *Nature*, 425(6956), 390–393.
- Shapiro, J. L. (1998). Does data-model co-evolution improve generalization performance of evolving learners? In Eiben, A. E., Bäck, T., Schoenauer, M., & Schwefel, H.-P. (Eds.), *Parallel Problem Solving from Nature V*, pp. 540–549. Springer-Verlag, Berlin, Germany.
- Shine, W., & Eick, C. (1997). Visualizing the evolution of genetic algorithm search processes. In *IEEE International Conference on Evolutionary Computation*, pp. 367–372. IEEE Press, Piscataway NJ.
- Shipman, R., Shackleton, M., Ebner, M., & Watson, R. (2000). Neutral search spaces for artificial evolution: A lesson from life. In Bedau, M. A., McCaskill, J. S., Packard, N. H., & Rasmussen, S. (Eds.), *Artificial Life VII*, pp. 162–169. MIT Press, Cambridge MA.
- Sih, A. (1998). Game theory and predator-prey response races. In Dugatkin, L. A., & Reeve, H. K. (Eds.), *Game Theory and Animal Behaviour*, pp. 221–238. Oxford University Press.
- Sims, K. (1991). Artificial evolution for computer graphics. *Computer Graphics*, 25(4), 319–328.
- Sims, K. (1994a). Evolving 3D morphology and behavior by competition. In Brooks & Maes (1994), pp. 28–39.
- Sims, K. (1994b). Evolving virtual creatures. In Glassner, A. (Ed.), *International Conference on Computer Graphics and Interactive Techniques*, pp. 15–22. ACM Press, New York NY.
- Sinervo, B., & Lively, C. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380, 240–243.

- Sklar, E., & Pollack, J. B. (2000). An evolutionary approach to guiding students in an educational game. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H. L., & Wilson, S. W. (Eds.), *Sixth International Conference on Simulation of Adaptive Behavior*, pp. 529–538. MIT Press, Cambridge MA.
- Slatkin, M., & Maynard Smith, J. (1979). Models of coevolution. *Quarterly Review of Biology*, *54*, 233–263.
- Slatkin, M. (1983). Genetic background. In Futuyma & Slatkin (1983), pp. 14–32.
- Smith, R. E., Forrest, S., & Perelson, A. S. (1993). Searching for diverse, cooperative populations with genetic algorithms. *Evolutionary Computation*, *1*(2), 127–149.
- Smith, T. M. C., Husbands, P., Layzell, P., & O’Shea, M. (2002). Fitness landscapes and evolvability. *Evolutionary Computation*, *10*(1), 1–34.
- Sober, E. (1984). *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. The University of Chicago Press, Chicago IL.
- Spears, W. (1994). Visualizing genetic algorithms. Technical report AIC-94-055, AI Center, Naval Research Laboratory, Washington DC.
- Spencer, H. (1864). *The Principles of Biology*, Vol. One. Williams and Norgate, London.
- Spencer, L., Goodman, E. D., Wu, A., Langdon, W., H., Gen, M., Sen, S., Dorigo, M., Pezeshk, S., Garzon, M. H., & Burke, E. (Eds.). (2001). *Genetic and Evolutionary Computation Conference*. Morgan Kaufmann, San Francisco CA.
- Stanley, K. O., & Miikkulainen, R. (2002). The dominance tournament method of monitoring progress in coevolution. In Barry (2002), pp. 242–248.
- Stanley, S. M., Van Valkenburgh, B., & Steneck, R. S. (1983). Coevolution and the fossil record. In Futuyma & Slatkin (1983), pp. 328–349.
- Tesauro, G. (1992). Practical issues in temporal difference learning. *Machine Learning*, *8*(3/4), 257–277.
- Thompson, J. N. (1982). *Interaction and Coevolution*. John Wiley, New York NY.
- Thompson, J. N. (1994). *The Coevolutionary Process*. University of Chicago Press, Chicago IL.
- van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, *1*, 1–30.



- von Neumann, J., & Morgenstern, O. (1953). *Theory of Games and Economic Behaviour*. Princeton University Press, Princeton NJ.
- Walker, S. (2000). *Detecting cyclic co-evolutionary trajectories in co-evolving genetic algorithms*. Unpublished B.Sc. dissertation, School of Computing, University of Leeds, UK.
- Wallace, A. R. (1858). On the tendency of varieties to depart indefinitely from the original type. *Proceedings of the Linnean Society (Zoology)*, 3, 53–62.
- Wanniski, J. (1978a). Taxes, revenues, and the ‘Laffer Curve’. *The Wall Street Journal*, June. Adaptation of Wanniski (1978b), Chapter 6.
- Wanniski, J. (1978b). *The Way the World Works: How Economies Fail—and Succeed*. Basic Books, New York NY.
- Watson, J. D., & Crick, F. H. C. (1953). Molecular structure of nuclear acids: A structure for deoxyribose nucleic acid. *Nature*, 171, 737–738.
- Watson, R. A., & Pollack, J. B. (2000). Symbiotic combination as an alternative to sexual recombination in genetic algorithms. In Schoenauer, M., Deb, K., Rudolph, G., Yao, X., Lutton, E., Merelo, J. J., & Schwefel, H. (Eds.), *Parallel Problem Solving from Nature VI*. Springer-Verlag, Berlin, Germany.
- Watson, R. A., & Pollack, J. B. (2001). Coevolutionary dynamics in a minimal substrate. In Spencer et al. (2001), pp. 702–709.
- Waxman, D., & Peck, J. R. (1999). Sex and adaptation in a changing environment. *Genetics*, 153(2), 1041–1053.
- Werner, G., & Dyer, M. (1992). Evolution of communication in artificial organisms. In Langton et al. (1992), pp. 659–687.
- Wiegand, R. P., Liles, W. C., & De Jong, K. A. (2001). An empirical analysis of collaboration methods in cooperative coevolutionary algorithms. In Spencer et al. (2001), pp. 1235–1245.
- Wiles, J., & Tonkes, B. (2002). Visualisation of hierarchical cost surfaces for evolutionary computation. In Fogel, D. (Ed.), *Congress on Evolutionary Computation*, pp. 157–162. IEEE Press, Piscataway NJ.

- Wolpert, D. H., & Macready, W. G. (1995). No free lunch theorems for search. Tech. rep. SFI-TR-95-02-010, Santa Fe Institute, Santa Fe NM.
- Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. In Jones, D. F. (Ed.), *International Proceedings of the Sixth International Congress on Genetics*, Vol. 1, pp. 356–366.
- Wright, S. (1967). Surfaces of selective value. *Proceedings of the National Academy of Science*, 58, 165–179.
- Wu, A. S., De Jong, K. A., Burke, D. S., Ramsey, C. L., & Grefenstette, J. J. (1999). VIS: A genetic algorithm visualization tool. In Wu, A. S. (Ed.), *Evolutionary Computation Visualization Workshop, Genetic and Evolutionary Computation Conference*. Morgan Kaufmann, San Francisco CA.
- Wynne-Edwards, V. C. (1962). *Animal dispersion in relation to social behaviour*. Oliver and Boyd, Edinburgh, UK.
- Yao, X., Liu, Y., & Darwen, P. (1996). How to make best use of evolutionary learning. In Stocker, R., Jelinek, H., & Durnota, B. (Eds.), *Complex Systems: From Local Interactions to Global Phenomena*, pp. 229–242. IOS Press, Amsterdam, Holland.
- Yao, X. (1997). Automatic acquisition of strategies by co-evolutionary learning. In Verma, B., & Yao, X. (Eds.), *International Conference on Computational Intelligence and Multimedia Applications*, pp. 23–29. Watson Ferguson & Company, Brisbane NSW.
- Zamudio, K. R., & Sinervo, B. (2000). Polygyny, mate-guarding, and posthumous fertilization as alternative male mating strategies. *Proceedings of the National Academy of Sciences USA*, 97(26), 14427–14432.