

Homeostatic Adaptive Networks

by

Hywel Thomas Parker Williams

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



**The University of Leeds
School of Computing**

June 2006

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.

Abstract

Homeostasis is constancy in the face of perturbation. The concept was originally developed to describe the fixed internal environment of an organism and this descriptive view of homeostasis has been prevalent in the literature. However, homeostasis can also be seen as the dynamic *process* of self-regulation and as such it is an organising principle by which systems adapt their behaviour over time. In this thesis we adopt this causal view of homeostasis and develop a theory of homeostatic adaptive systems.

We study homeostatic adaptive networks by looking at specific examples of homeostatic systems: the Homeostat, homeostatic plasticity in neural networks, and homeostatic regulation of the environment by the biota. Investigation of these case studies forms the basis for the development of a generalised theory of homeostatic adaptive systems.

The Homeostat was an electromechanical device designed by W. Ross Ashby to demonstrate the principle of ultrastability, where the stability of a system requires homeostasis of essential variables. Ashby put forward a theory of mammalian learning as a process of homeostatic adaptation that was based on the idea of the ultrastable system. Here we develop a simulated Homeostat and explore its properties as a homeostatic adaptive system, looking at its ultrastable nature and its ability to adapt to external perturbations.

The second case study, neural homeostasis, has recently been a topic of much interest in the neurosciences, with new data being presented concerning the existence and functioning of a variety of mechanisms by which neural activity is regulated. Homeostatic plastic mechanisms prevent long term quiescence or hyper-excitation in biological neurons and this suggests that such mechanisms may be used to solve the problem of node saturation in artificial neural networks. Here we develop homeostatic plastic mechanisms for use in continuous-time recurrent neural networks, a kind of network often used in evolutionary robotics, and study the effect of these mechanisms on network behaviour. Node saturation effects can make these networks difficult to evolve as robot controllers and we also look at the effect of homeostatic plasticity on evolvability.

The third case study is the evolution of homeostatic regulation of the physical environment by the biota. The Gaia theory states that life regulates the entire biosphere to conditions suitable for life, but the general concept of biological regulation of the environment is applicable on a variety of scales. However, there are major theoretical issues concerning the compatibility of environmental regulation with evolutionary theory. Here we develop a modified version of the Daisyworld model and use it to determine the compatibility of global regulation with individual selection. We show that regulation in Daisyworld depends on several key assumptions and fails if these assumptions are removed. We develop the Flask model, in which environmental regulation by microbial communities evolves as

a result of multi-level selection, in order to show how regulation can occur when the core assumptions of Daisyworld are relaxed.

At the end of the thesis we try to draw some general conclusions concerning homeostatic adaptive systems. We consider the adaptive and homeostatic properties of each of the case study systems, and then generalise from these to give some principles of homeostatic adaptation. Our analysis shows that perturbations to a system can be classified in terms of their effect on homeostasis, and that the ability of a system to adapt to a perturbation and maintain homeostasis depends on the variety of responses it can produce. We argue that parameter change caused by a loss of homeostasis causes ‘organisation death’ in a homeostatic adaptive system, where the system does not survive in its current form. This suggests a view of learning and evolution of organisms as second order homeostatic adaptive processes.

Acknowledgements

This thesis is the culmination of several years work and the author would like to thank everyone who has contributed to the formation of the ideas within it, at any point and in whatever way. Certain parties deserve a special mention: the anonymous reviewers who have commented on those parts of the thesis which have appeared in the literature, anyone who has ever commented on my work or with whom I have discussed my ideas at conferences, the people (especially Dan Franks and Jason Noble) who have proof-read sections of the thesis and greatly improved its quality. Any remaining errors or omissions are entirely my own responsibility.

The Biosystems Group at the University of Leeds has been a stimulating and supportive environment in which to conduct research. The weekly reading group meetings have been variously interesting, enlightening and infuriating, but have never failed to encourage fresh thinking and new ideas. Every individual in the group has contributed something to my work during my time there, not least a friendly atmosphere — thanks guys.

A few individuals have made an outstanding impact on my work. Chris Buckley has provided a constant source of enthusiasm and rigour that has forced me to think clearly and do my sums. Also, my supervisor Jason Noble has been steadily supportive and has listened patiently as my vague ideas took shape. His mix of incisive critical comment and laid back approach has sharpened my ideas while allowing me to go my own way; I am glad to have been his student. In terms of ideas, the big three have got to be Charles Darwin, William Ross Ashby, and James Lovelock.

On a personal level, this thesis is dedicated to the people who have made it possible: my parents Gill and Steve, for their steady belief and support, Chris, without whom I would never have begun, and Abby, for her understanding and for reminding me that there is more to life than work.

Declarations

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

Williams, H. and Noble, J. “Evolution and the regulation of environmental variables”. In Capcarrere, M.S., Freitas, A.A., Bentley, P.J., Johnson, C.G., and Timmis, J., editors, *Proceedings of the 8th European Conference on Artificial Life (ECAL 2005)*, pages 332-341, Berlin Heidelberg, 2005. Springer-Verlag.

Williams, H. and Noble, J. “Homeostatic plasticity improves signal propagation in continuous-time recurrent neural networks”. To appear in a special issue of *BioSystems*.

Williams, H. “Homeostatic plasticity improves continuous-time recurrent neural networks as a behavioural substrate”. *Proceedings of 3rd International Symposium on Adaptive Motion in Animals and Machines, Ilmenau, Germany (AMAM 2005)*, 2005.

Williams, H. “Homeostatic plasticity in recurrent neural networks”. In Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J-A., editors, *From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior*, 2004, Cambridge, MA. MIT Press.

Contents

I	Introduction	1
1	Homeostasis as an organising principle	2
1.1	Heating the home	2
1.2	Overview	4
1.3	Homeostasis	4
1.3.1	A brief history of homeostasis	4
1.3.2	Examples of homeostasis	5
1.3.2.1	Physiological homeostasis	5
1.3.2.2	Ecological homeostasis	6
1.3.2.3	Genetic homeostasis	7
1.3.3	Homeostasis as an organising principle	8
1.3.4	Definitions	9
1.3.4.1	Homeostasis	9
1.3.4.2	Homeostatic adaptation	11
1.3.4.3	Multiple homeostatic feedbacks	12
1.3.4.4	Second order homeostasis	12
1.3.5	Examples from a minimal homeostatic system	13
1.3.5.1	An elastic homeostatic adaptive system	14
1.3.5.2	A plastic homeostatic adaptive system	14
1.3.5.3	A plastic homeorhetic adaptive system	15
1.3.5.4	A homeostatic adaptive system with multiple homeo- static feedback loops	15
1.3.5.5	A second order homeostatic adaptive system	16
1.4	Thesis outline	16
1.4.1	Research aims	16
1.4.2	Case studies	17
1.4.2.1	Ashby's Homeostat	17
1.4.2.2	Neural homeostasis	17

1.4.2.3	Environmental regulation by the biota	18
1.4.3	Methodology	18
1.4.4	Thesis plan	20
II	The Homeostat	21
2	The Homeostat	22
2.1	Overview	22
2.2	Background	22
2.2.1	Ultrastability	23
2.2.2	The Homeostat	24
2.3	Method	25
2.4	Results	27
2.4.1	Reaching stability	27
2.4.2	Time to convergence	30
2.4.3	Perturbations	34
2.4.4	A control problem	34
2.5	Discussion	37
III	Homeostatic Plasticity in Neural Networks	42
3	Background to Part III	43
3.1	Overview	43
3.2	Neural and evolutionary robotics	44
3.3	Evolving robot controllers	45
3.3.1	Genetic algorithms	45
3.3.1.1	Encoding	47
3.3.1.2	Crossover	47
3.3.1.3	Mutation	47
3.3.1.4	Selection	48
3.3.1.5	Fitness function	49
3.3.2	Using GAs for evolutionary robotics	50
3.4	Continuous-time recurrent neural networks	50
3.4.1	Mathematical formulation of CTRNNs	51
3.4.2	Evolving CTRNNs	52
3.4.3	Improving evolvability	52

3.4.4	The saturation problem	53
3.5	Homeostatic plasticity in biological nervous systems	56
3.5.1	Plasticity of intrinsic excitability of neurons	57
3.5.2	Synaptic scaling in the central nervous system	58
3.5.3	Synaptic homeostasis at the neuromuscular junction	59
3.5.4	Neurotrophins and regulation of activity in cortical networks	59
3.5.5	Activity-dependent regulation of synapse number	59
3.6	Homeostatic plasticity in artificial nervous systems	60
3.6.1	Di Paolo (2000) ‘Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions’	60
3.6.2	Di Paolo (2003) ‘Organismically-inspired robotics: homeostatic adaptation and teleology beyond the closed sensorimotor loop’	62
3.6.3	Di Paolo (2002) ‘Evolving robust robots using fast homeostatic oscillators’	63
3.6.4	Mathayomchan & Beer (2002) ‘Centre-crossing recurrent neural networks for the evolution of rhythmic behaviour’	64
3.6.5	Hoinville & Henaff (2004) ‘Evolving plastic neural controllers stabilized by homeostatic mechanisms for adaptation to a perturbation’	65
3.7	Can simple simulations inform neuroscience?	67
3.8	Aims, objectives, questions	69
4	Can homeostatic plasticity be implemented in CTRNNs?	71
4.1	Overview	71
4.2	What types of mechanism should be implemented?	71
4.3	Functional effects of homeostatic plastic mechanisms	72
4.3.1	Synaptic scaling	73
4.3.2	Plasticity of intrinsic excitability	73
4.4	Mathematical formulation of homeostatic plastic mechanisms	73
4.4.1	Plastic facilitation	74
4.4.2	Synaptic scaling	74
4.4.3	Adaptive bias	75
4.5	Can homeostatic plasticity be applied to other kinds of network?	76
4.6	Neuroscientific relevance	76
5	Does homeostatic plasticity prevent node saturation?	77
5.1	Overview	77

5.2	Nature of homeostatic constraint	77
5.3	Satisfaction of homeostatic constraint in a single node	79
5.4	Equilibrium points for the single-node system	81
5.5	Behaviour of the single-node system in the unistable and bistable regimes	82
5.6	Effect of homeostatic plasticity on equilibria	85
5.7	Constraint satisfaction in a network	86
5.8	Implications for neuroscience	88
6	What effect does homeostatic plasticity have on network behaviour?	89
6.1	Overview	89
6.2	Methodology	89
6.2.1	Plasticity, learning, and progress	90
6.2.2	Measuring the effects of plasticity	90
6.2.3	Adiabatic approximation	91
6.2.4	Lifetime plasticity versus developmental period	92
6.2.5	Measurement from random ensembles	92
6.2.6	Metrics for measuring networks	93
6.2.6.1	Useful qualities for behavioural substrates	93
6.2.6.2	Metrics	94
6.2.7	Levels of analysis	94
6.3	Method	95
6.3.1	Photo-sensitive robot simulator	95
6.3.1.1	Kinematics	95
6.3.1.2	Light sensors	96
6.3.1.3	Control network	97
6.3.1.4	Environment	97
6.3.2	Metrics	97
6.3.2.1	Node sensitivity	97
6.3.2.2	Signal propagation	98
6.3.2.3	Oscillations	98
6.4	Node-level effects of homeostatic plasticity	100
6.4.1	Slow oscillations	100
6.4.2	Increased sensitivity	101
6.5	Network-level effects of homeostatic plasticity	102
6.5.1	Additional behavioural timescales	102
6.5.2	Transient dynamics	102
6.5.3	Improved signal propagation	104

6.5.3.1	Fully connected networks	105
6.5.3.2	Randomly connected networks	105
6.5.3.3	Feed-forward networks	106
6.5.4	Greater likelihood of oscillations	108
6.6	Agent-level effects of homeostatic plasticity: anecdotal evidence	109
6.6.1	Robots controlled by homeostatic plastic CTRNNs	110
6.6.2	Robots controlled by non-plastic CTRNNs (before/after plasticity)	111
6.7	Summary of results	113
6.8	Implications for robotics	113
6.9	Implications for neuroscience	115
7	Does homeostatic plasticity make CTRNNs more evolvable?	117
7.1	Overview	117
7.2	Homeostatic plastic CTRNNs as an evolutionary substrate	117
7.3	Measuring evolvability	118
7.4	Method	119
7.4.1	Agent	120
7.4.1.1	Control network	120
7.4.1.2	Sensors	121
7.4.1.3	Motors	121
7.4.2	Falling shapes	121
7.4.3	Genetic algorithm	122
7.4.3.1	Encoding	122
7.4.3.2	Selection	122
7.4.3.3	Mutation / crossover	123
7.4.3.4	Fitness function 1: Ball-catching task	123
7.4.3.5	Fitness function 2: Discrimination task	124
7.5	Results	125
7.5.1	Experiment 1: Comparative evolvability of CTRNNs with different homeostatic plasticity schemes	125
7.5.2	Experiment 2: The effect of a developmental period	131
7.5.3	Experiment 3: Comparison of pre-conditioned networks with centre-crossing networks	136
7.5.4	Experiment 4: Evolvability of larger non-plastic CTRNNs	139
7.6	Summary	142
7.7	Implications for robotics	143
7.8	Implications for neuroscience	145

8	Discussion of Part III	146
8.1	Overview	146
8.2	Summary of results	146
8.2.1	Chapter 3: Background	146
8.2.2	Chapter 4: Can homeostatic plasticity be implemented in CTRNNs?	147
8.2.3	Chapter 5: Does homeostatic plasticity prevent node saturation?	148
8.2.4	Chapter 6: What effect does homeostatic plasticity have on network behaviour?	148
8.2.5	Chapter 7: Does homeostatic plasticity make CTRNNs more evolvable?	149
8.3	Implications for robotics	150
8.4	Implications for neuroscience	152
8.5	Review and Future Work	154
8.6	Conclusions	155
IV	Homeostatic Regulation of the Environment by the Biota	157
9	Background to Part IV	158
9.1	Overview	158
9.2	Gaia theory	160
9.2.1	Precursors: Vernadsky	160
9.2.2	Early development	160
9.2.3	Criticisms	161
9.2.4	Daisyworld	162
9.2.5	Multiple Gaia hypotheses	162
9.2.6	Gaia vs Evolutionary biology	163
9.2.6.1	Early criticisms	164
9.2.6.2	The case <i>for</i> the evolution of regulation	164
9.2.6.3	The case <i>against</i> the evolution of regulation	165
9.2.6.4	By-product Gaia	166
9.2.7	Gaia - an organism?	166
9.2.8	Current status of the Gaia theory	168
9.3	Models of Gaia	169
9.3.1	Daisyworld	169
9.3.1.1	The original Daisyworld model	170
9.3.1.2	Ecology in Daisyworld	174
9.3.1.3	Chaos in Daisyworld	175

9.3.1.4	Spatial Daisyworlds	175
9.3.1.5	Evolution in Daisyworld	176
9.3.1.6	Cut-down Daisyworld	177
9.3.2	Other Gaia models	178
9.3.2.1	The Guild model	178
9.4	Aims, objectives, questions	180
10	Evolution in Daisyworld	183
10.1	Overview	183
10.2	Rationale	183
10.3	Model	184
10.3.1	Daisies	184
10.3.2	Seeding	185
10.3.3	Colonisation	186
10.3.4	Death	186
10.3.5	Calculation of patch temperature	186
10.3.6	Cellular automata update	187
10.4	Results 1: Replication of existing Daisyworld results	187
10.5	Results 2: Constraints on evolution and their implications for environmental regulation	190
10.6	Results 3: Relaxing core assumptions	193
10.7	Conclusion	196
11	The Flask model	198
11.1	Overview	198
11.2	Moving on from Daisyworld	199
11.3	Biological basis for the Flask model	201
11.4	Differences from existing models	202
11.5	The Flask model: A conceptual sketch	203
11.6	Basic assumptions and points to prove	209
11.6.1	Universal preferred environment?	209
11.6.2	Environmental fit correlates with biomass?	210
11.6.3	Heritability at flask level?	210
11.6.4	Spatial structure leads to group selection?	210
11.7	Conclusion	210

12 Response to ecosystem selection in the Flask model	212
12.1 Overview	212
12.2 Background	213
12.3 Method	214
12.3.1 Flasks	215
12.3.2 Microbes	216
12.3.2.1 Genotype	216
12.3.2.2 Metabolism	216
12.3.2.3 Guilds and species	218
12.3.2.4 Reproduction	218
12.3.2.5 Death	219
12.3.3 Update scheme	219
12.3.4 Artificial selection experiments	219
12.4 Results	220
12.4.1 Ecosystem dynamics	220
12.4.2 Response to selection	221
12.4.3 Time-dependence of the response to selection	223
12.5 Discussion	224
12.6 Conclusion	225
13 Discussion of Part IV	226
13.1 Overview	226
13.2 Summary of results	226
13.2.1 Chapter 9: Background	226
13.2.2 Chapter 10: Evolution in Daisyworld	227
13.2.3 Chapter 11: The Flask model	229
13.2.4 Chapter 12: Response to ecosystem selection in the Flask model	229
13.3 Implications for Gaia theory	230
13.4 Implications for biology	232
13.5 Review and future work	233
13.6 Conclusions	234
V Discussion	235
14 Towards a general theory of homeostatic adaptive systems	236
14.1 Overview	236
14.2 Definitions	237

14.3	Case study 1: The Homeostat	237
14.4	Case study 2: Homeostatic plasticity in neural networks	239
14.5	Case study 3: Homeostatic regulation of the environment by the biota	241
14.5.1	Daisyworld	241
14.5.1.1	Two-species Daisyworld	241
14.5.1.2	Evolutionary Daisyworld	242
14.5.2	Flask model	245
14.5.3	The Earth system	245
14.6	Discussion	246
14.6.1	Classification of perturbations	246
14.6.2	Requisite variety	247
14.6.3	Relation between first and second order systems	248
14.6.4	Parameter selection mechanism	250
14.6.5	Organisation death and essential variables	250
14.6.6	Learning as second order homeostasis	251
14.6.7	Evolution as second order homeostasis	252
14.7	Conclusion	253
15	Conclusion	255
15.1	Overview	255
15.2	Part I: Introduction	255
15.3	Part II: The Homeostat	256
15.4	Part III: Homeostatic plasticity in neural networks	257
15.5	Part IV: Homeostatic regulation of the environment by the biota	259
15.6	PartV: Discussion	262
15.7	Conclusion	262
	Bibliography	264

List of Figures

1.1	Physiological homeostasis.	6
1.2	Thermoregulation.	6
1.3	Blood glucose homeostasis.	7
1.4	Homeostasis.	10
1.5	Multiple homeostatic feedbacks.	12
1.6	Second order homeostasis.	13
1.7	Minimal homeostatic system.	14
2.1	Ultrastability.	24
2.2	Ashby's Homeostat.	25
2.3	Schematic of a single Homeostat unit.	26
2.4	Example Homeostat unit transfer function.	26
2.5	Convergence in a 4-unit Homeostat with transfer function adaptation.	28
2.6	Convergence/non-convergence with adaptation of connection strengths.	29
2.7	Probability of homeostasis with adapting connection strengths.	31
2.8	Probability of homeostasis with adapting transfer functions.	32
2.9	4-unit Homeostat where random input is applied to one unit.	35
2.10	4-unit Homeostat with perturbation applied to all units.	36
2.11	4-unit Homeostat with designated output unit.	38
3.1	Schematic of a single node in a CTRNN network.	51
3.2	Sigmoidal transfer function.	54
3.3	Equilibrium firing rates / firing rate responses to a sinusoidal input for 3 example nodes.	55
3.4	Homeostatic plasticity of intrinsic excitability of neurons.	57
3.5	Homeostatic plasticity by synaptic scaling.	58
4.1	Schematic of effects of homeostatic plasticity mechanisms.	72
4.2	Plastic facilitation.	74

5.1	Constraint-satisfying region of parameter space.	79
5.2	Schematic of a single-node CTRNN.	79
5.3	Homeostatic plasticity moves nodes to constraint-satisfying region.	80
5.4	Single-node equilibria plotted for 3 different nodes.	82
5.5	Equilibria for unistable and bistable nodes.	83
5.6	Vector field plots for unistable and bistable nodes.	84
5.7	Equilibria for unistable and bistable nodes showing target range.	85
5.8	Plasticity makes a bistable neuron unistable and satisfies constraint.	87
5.9	Plasticity leaves node bistable and fails to satisfy constraint.	87
6.1	Diagram of photosensitive robot.	96
6.2	Different network topologies used to test signal propagation.	99
6.3	Node sensitivity rises as homeostatic plasticity is applied.	101
6.4	Signal propagation in fully connected CTRNNs.	106
6.5	Signal propagation in CTRNNs based on Erdős-Renyi random graphs.	107
6.6	Signal propagation in feedforward networks.	108
6.7	Proportion of initial conditions that lead to oscillations.	110
6.8	Motion plots of robots controlled by homeostatic plastic CTRNNs.	112
6.9	Behaviour examples	114
7.1	Ray sensor agent.	120
7.2	Mean evolutionary performance on ball-catching task.	126
7.3	Individual run performances on ball-catching task.	127
7.4	Mean evolutionary performance on discrimination task.	128
7.5	Individual run performances on discrimination task.	129
7.6	Effect of development on plastic network on ball-catching task.	132
7.7	Effect of development on plastic network on discrimination task.	133
7.8	Effect of development on non-plastic network on ball-catching task.	134
7.9	Effect of development on non-plastic network on discrimination task.	135
7.10	Comparison with centre-crossing networks on ball-catching task.	137
7.11	Comparison with centre-crossing networks on discrimination task.	138
7.12	Effect of network size on ball-catching task.	140
7.13	Effect of network size on discrimination task.	141
9.1	Results from original Daisyworld.	173
9.2	Results from a Guild simulation run.	181
10.1	Example daisy growth function.	185
10.2	Results with pure black and pure white daisy species only.	188

10.3	Results when daisy albedo can mutate freely.	188
10.4	Results when albedo and preferred temperature mutate freely.	189
10.5	Constrained growth function.	189
10.6	Results when maximum growth rate is constrained.	190
10.7	Results with two possible preferred growth temperatures.	191
10.8	Results with two overlapping growth ranges (increasing luminosity).	192
10.9	Results with two overlapping growth ranges (decreasing luminosity).	192
10.10	Results with slow mutation of preferred growth temperature.	192
10.11	Results with no local buffer against environment.	195
10.12	Results with random selection.	196
11.1	Every organism has an environment.	203
11.2	Schematic of a single flask.	204
11.3	Metabolism is affected by the environment.	205
11.4	Colonisation and invasion are possible in the multi-flask model.	207
12.1	Ecological dynamics of a single flask ecosystem.	221
12.2	Mean fitness over generational time.	222
12.3	Response to selection against generation duration.	223
14.1	Homeostatic feedbacks in original Homeostat.	238
14.2	Homeostatic feedbacks in simulated Homeostat.	239
14.3	Homeostatic feedbacks in homeostatic plastic CTRNNs.	240
14.4	Homeostatic feedbacks in two-species Daisyworld.	242
14.5	Homeostatic feedbacks in evolutionary Daisyworld.	243
14.6	Organisation death in the Homeostat.	251

List of Tables

6.1	Convergence to attractor: Non-plastic CTRNNs.	103
6.2	Convergence to attractor: HP-CTRNNs with target range [0.05,0.95]. . .	103
6.3	Convergence to attractor: HP-CTRNNs with target range [0.25,0.75]. . .	103
14.1	Homeostatic outcomes of example system.	248
14.2	Homeostatic outcomes of second order system.	249

Part I

Introduction

Chapter 1

Homeostasis as an organising principle

1.1 Heating the home

The occupants of a house with no mechanisms for altering the internal temperature are subject to the whim of the elements. When the sun shines in summer they will be too hot, and when the wind blows in winter they will be too cold. Installing central heating and air conditioning allows the occupants to *regulate* the temperature of the house by making adjustments to the internal temperature when external conditions demand it. By switching on the heating when it is too cold, or the air conditioning when it is too hot, the occupants can keep the house close to their desired temperature.

To avoid having to repeatedly make manual adjustments to the temperature, it might be useful to install some form of automatic device that will monitor the temperature and make adjustments as necessary. Such a device, together with the heating and air conditioning systems, would make the house *homeostatic*. Temperature within the house would be controlled by a process of dynamic self-regulation. The occupants of one particular house narrow down their options to a choice of two devices for the automatic temperature monitoring and adjustment role. One device is a thermostat, and the other a Trial-o-matic™. Both devices connect to the heating and air conditioning systems to make the house homeostatic, but each has different strengths and weaknesses.

The thermostat continuously monitors room temperature and detects any deviation from the desired level. If temperature drops, the thermostat switches on the heating; if temperature rises, the thermostat activates the air conditioning. The major selling point

of the thermostat is its efficiency; when correctly wired it always picks the appropriate mechanism to counter-act any deviation from the optimal temperature. However, the thermostat has some limitations that arise from its fixed mode of operation. Suppose that the thermostat is wired up incorrectly to the heating and air conditioning systems, so that their connections are switched. Now when the house is too cold, the thermostat will activate the air conditioning, thus making the house colder; when the house is too hot, the thermostat will switch on the heating and make it hotter still. The regulatory system is broken, and there is no way for the thermostat to remedy the situation.

The Trial-o-matic is a different form of controller to the thermostat. It monitors the temperature of the house and responds to any deviation from the desired level by activating either the heating or the air conditioning system, but its choice of which system to use depends on its history. When newly installed, the Trial-o-matic responds to every temperature deviation by activating at random either the heating or the air conditioning, and then waiting for a period to see if this action has had the desired effect. If the temperature has returned to the desired level, the Trial-o-matic switches the system off and records which system it used to correct that sort of deviation. Alternatively, if the temperature is still not at the desired level, the Trial-o-matic repeats the random selection process until eventually the deviation is corrected. Over time the list of known deviation-response pairs grows, so that the Trial-o-matic is often faced with temperature changes it has experienced before; when this occurs it consults its memory to find the correct response. In this way, the Trial-o-matic slowly ‘learns’ the responses needed to regulate the temperature of the house and improves its efficiency.

While the thermostat is efficient but vulnerable to mistakes in its wiring, the Trial-o-matic is initially inefficient but able to adapt to different wiring schemes. Both the thermostat and the Trial-o-matic act to maintain homeostasis in room temperature, but do so in different ways. Each method has its strengths and weaknesses, which result from the different mechanisms they use to adapt to perturbations.

In reality, central heating in the home is a well understood problem and it is unlikely that there would ever be a need for a reconfigurable controller such as the Trial-o-matic. Thermostats are reliable and effective, and there is no need for an alternative method of temperature regulation. However, homeostatic regulation is not always a simple process, easily implemented with a thermostat. For example, when designing systems that are more complex and less well understood than central heating, it may not be possible to specify in advance the necessary responses needed to regulate a variable. This kind of situation may require a more sophisticated regulatory mechanism that can adapt its response. Also, when seeking to understand homeostatic systems in biology we can observe that evolved biological systems are likely to be messy and contingent on their evolution-

ary history, making regulation in these systems inefficient and dependent on mechanisms that did not originally evolve for that purpose. These observations mean that homeostatic adaptive systems are not necessarily easy to understand and simple to design. There is a need for a better understanding of how homeostasis can be achieved in complex systems.

1.2 Overview

This thesis looks at homeostasis as a cybernetic organising principle. It considers homeostasis not as a purely descriptive property, but as a process of dynamic self-regulation that shapes the behaviour of a system. We will study the different forms of adaptive mechanism that can lead to homeostasis, their similarities and their differences, by looking at three example homeostatic adaptive systems: the Homeostat, neural homeostasis, and environmental regulation by the biota. We will then build on these case studies to try and form a general theoretical understanding of the properties of homeostatic adaptive networks.

This chapter first of all gives some historical background to the concept of homeostasis and offers a few examples of homeostatic systems. Next some concepts related to homeostasis are defined, with reference to a toy model of a minimal homeostatic system. This is followed by a statement of the research aims of the thesis, together with a brief introduction to the case study systems that provide the bulk of the thesis material. The chapter concludes with a few comments on methodology and on the layout of the thesis.

1.3 Homeostasis

1.3.1 A brief history of homeostasis

The term *homeostasis* was originally coined in 1932 by Walter Cannon [26], but the concept owes much to the earlier work of Claude Bernard (1813-1878), a French scientist who was perhaps the first to recognise that organisms maintain a near-constant internal environment despite changing external influences. Bernard famously wrote that “La fixité du milieu intérieur est la condition de la vie libre”, which translates as “The fixity of the internal environment is the condition for free life” [20]. Cannon, an American physiologist, derived the term *homeostasis* from the Greek *homo* (same, like) and *stasis* (to stand, posture). Cannon used the term to describe the properties of stability and constancy that he observed in living organisms. Cannon studied the ability of living organisms to regulate their internal environment, detailing many of the mechanisms by which various elements

in the human body are regulated in his book ‘The Wisdom of the Body’ [26]. These included regulation of body temperature through the skin, regulation of blood glucose by the liver and pancreas, regulation of water content by the kidneys, and the role of stimuli such as hunger and thirst in ensuring the supply of necessary materials.

Homeostasis remained a physiological term until William Ross Ashby, a British cybernetician, proposed an alternative view of homeostasis that implicated the regulation of ‘essential variables’ in the generation of behaviour [7]. Ashby developed the linked concepts of *ultrastability* and *homeostatic adaptation* as an explanation of learning in organisms and machines (these concepts will be explained later). Ashby’s ideas were theoretically advanced, but their development was hindered by the resources available to him; only recently with the advent of easily available computational power are his ideas beginning to be explored [46].

It is Ashby’s concept of homeostasis as a dynamic adaptive process that we are most concerned with here, but it is interesting to note that homeostasis is a widely recognised phenomenon in a variety of fields. It is still seen primarily as a physiological quality, but homeostasis has been applied and studied in areas as diverse as ecology, control theory, psychology, neuroscience, and transport.

1.3.2 Examples of homeostasis

1.3.2.1 Physiological homeostasis

Homeostasis as a concept was originally developed to describe the internal constancy of the human body, and it is in this area that most examples can be found. Many aspects of the body are regulated homeostatically, including blood glucose level, body temperature, blood water content, pH, etc. Physiological (and other) homeostatic systems typically function in the manner shown in Figure 1.1. Receptors detect a change in the level of the target variable and trigger a compensatory action in effectors. There may be a variety of different receptors and effectors involved in the regulation of a single variable. For example, body temperature homeostasis involves detectors in the hypothalamus and in the skin, and several different effector mechanisms (Figure 1.2). It is common for a homeostatic system to use different mechanisms to move the target variable in opposing directions. Blood glucose level is controlled in the pancreas by two hormones, insulin and glucagon. Insulin decreases the level of glucose in the blood, while glucagon raises the level of glucose in the blood. Changes in blood glucose concentration are detected in the pancreas. If glucose level falls too low, then the pancreas secretes glucagon, while if glucose level rises too high, the pancreas secretes insulin (Figure 1.3). Physiological homeostasis can involve both behaviour and internal mechanisms. For example, the amount of water in

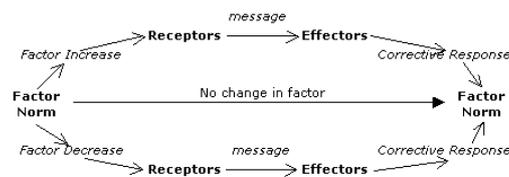


Figure 1.1: Physiological homeostasis. Deviation of a factor from the normal level is detected by receptors that trigger effectors to compensate and return the factor to its normal level. Diagram reproduced with permission from http://www.biology-online.org/4/1_physiological_homeostasis.htm, 24th January 2006.

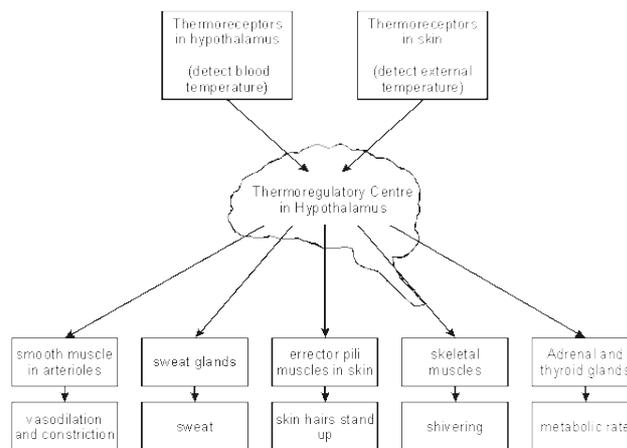


Figure 1.2: Thermoregulation involves a variety of different receptors and effectors. Changes in temperature are detected by receptors in the skin and hypothalamus, which trigger compensatory changes via effectors in skin and muscle. Diagram taken from <http://www.biologymad.com/Homeostasis/Homeostasis.htm>, 24th January 2006.

the blood is regulated via the hypothalamus both by triggering sensations of thirst and by controlling the reabsorption of water from urine.

1.3.2.2 Ecological homeostasis

The most famous example of ecological homeostasis is the Gaia theory [116], which postulates global regulation of the environment by the biota. Gaia theory states that the collective effect of the interactions between the biota and the physical environment will be such that the biosphere is homeostatically regulated in a range suitable for life. This controversial theory has slowly become more accepted by the scientific community, although uncertainties still exist, particularly concerning the compatibility of the Gaia hypothesis with evolutionary biology. Gaia theory will be covered in more detail in Part IV of this thesis.

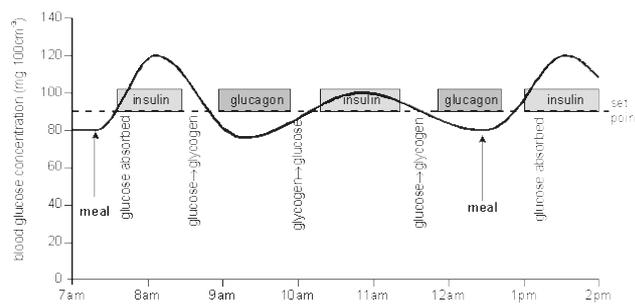


Figure 1.3: Blood glucose homeostasis uses different mechanisms to move the target variable in different directions: insulin and glucagon secreted in the pancreas counteract changes in blood glucose level by adjusting it downwards and upwards respectively. Diagram taken from <http://www.biologymad.com/Homeostasis/Homeostasis.htm>, 24th January 2006.

Another way in which homeostasis is thought to be expressed at an ecosystem level is in regulation of the population size of a species by density dependent factors [19]. These might include limited food supply, predation, disease, etc. If the population density of a species affects either its birth rate or death rate, then there will be some equilibrium value for the population size when birth rate is equal to death rate and the population size is stable. Below this equilibrium, the population will increase, while above it the population size will decrease. Thus density dependent population growth is a form of ecological homeostasis.

A more sophisticated notion of ecosystem homeostasis is put forward by Trojan [166], who considers ecological equilibrium as a form of homeostasis. Trojan notes that when an ecosystem reaches an equilibrium state, there will often be mechanisms by which this state is maintained in the face of disturbance. In Trojan's view, the key regulatory processes in an ecosystem are those protecting matter cycling and energy flow, primary production level, and system structure. The structure of the ecosystem is thus both the mechanism and part of the target for homeostatic regulation.

1.3.2.3 Genetic homeostasis

Lerner [112] developed a theory of genetic homeostasis, which he described as the tendency of Mendelian populations to maintain a constant genetic composition in the face of external pressure. Lerner noted that although selection operates on the level of individuals, a by-product of segregation is that gene pools tend to become integrated and maintain an optimal balance of gene frequencies at different loci. When subjected to pressure, such as that from artificial selection experiments (the area where most of Lerner's data came from), genetic homeostatic mechanisms acted to restore any frequencies which

have shifted from mean optimal values.

1.3.3 Homeostasis as an organising principle

Homeostasis is a widespread phenomenon, with many more examples than the few given above. There is considerable variety in both the targets and the mechanisms of regulation, and in the physical and temporal scales of different homeostatic systems. In this thesis we are interested in homeostasis as an organising principle, that is, as a force that can shape the dynamics and behaviour of a system. This view moves away from the more traditional view of homeostasis as simply a property displayed by certain types of system; we want to move away from this descriptive viewpoint towards a concept of homeostasis that gives it a causal role in the development of systems over time. In effect this switches from an ‘adjective’ definition of homeostasis to a definition that treats it as a verb.

Consider the alternative devices for regulating room temperature that were discussed in the anecdote given at the start of this chapter. Both the thermostat and the Trial-o-matic maintain homeostasis in temperature by triggering the activation of either the heating or the air conditioning systems when a deviation in temperature occurs. However, apart from this high level similarity, the two regulatory mechanisms are qualitatively different in nature. The thermostat has a fixed response, in that it will always respond in the same way to a given ambient temperature. Since this response has been designed to move temperature towards the desired target level, the thermostat is very efficient. However, it is also brittle, so that any error in its wiring will lead to a failure of regulation. The Trial-o-matic, by contrast, has a response that varies over time as it ‘learns’ to regulate the house, so that it might give a different response to the same temperature on two different occasions. Its design is such that its response should improve over time, as it records more effective temperature-response pairings. The trial-and-error method by which the Trial-o-matic acquires better responses means that it is likely to be inefficient early in its operational lifetime, but offers the advantage of robustness. Because it has no pre-ordained responses, it will adapt to being wrongly connected to its heating and cooling effector systems.

The thermostat can be seen as an elastic system, where changes in behaviour have no persistent effect. Its behaviour may change temporarily to accommodate a perturbation, but when the temperature is restored to the desired level the thermostat will return to a resting condition identical to its original state. The Trial-o-matic, however, is a plastic system, where changes can be persistent. The Trial-o-matic’s current behaviour can affect its future behaviour, so that its behaviour at any point in time is a function both of its current environmental stimulus and of its operational history.

Elastic homeostatic systems are well studied and include most of the physiological examples of homeostasis given above. This form of homeostatic system is the domain of control theory, where explanations often involve negative feedback loops. Plastic homeostatic systems are much less well studied. Plastic homeostatic systems change their behaviour semi-permanently or permanently depending on their precise formulation, and this means that their continual drive towards homeostasis is a force that shapes their organisation over time. The persistent changes to the system that occur as the system attempts to regulate its target variables have an effect on its subsequent behaviour that will last until over-written by some future changes.

It is plastic homeostatic systems that are the topic of this thesis. We will look at three examples of plastic homeostatic systems as case studies, and then look for any system features or properties that are general to all. The three case studies are the Homeostat, homeostatic plasticity in neural networks, and environmental regulation by the biota.

1.3.4 Definitions

Before we proceed, we must first of all attempt some more rigorous definitions of homeostasis and related concepts. While definitions make no difference to the behaviour of a system, precise terminology will be of use when trying to compare and contrast the different case study systems in the final stages of this study. Here we present some definitions of various forms of homeostasis, followed by some examples based on a minimal homeostatic system.

It is not straightforward to give a rigorous definition of homeostasis that clearly distinguishes it from related concepts such as negative feedback or system stability; the definitions given below represent our best attempt. There is overlap between these concepts, and there is a danger of lapsing into arbitrary semantics when trying to separate them. Homeostatic systems can display negative feedback, but need not do so. Homeostatic systems show stability, but the term homeostasis implies more about the nature of a system than just stability. In the interests of avoiding a lengthy linguistic discussion and retaining an effective terminology, we will here give a set of definitions that is pragmatically focussed on ease of use and clarity. At all points during this thesis, systems that are described as homeostatic will also be described mathematically and mechanistically, so that no confusion over function should arise.

1.3.4.1 Homeostasis

Homeostasis is the maintenance by some mechanism(s) of a variable within a target range, in opposition to forces which would otherwise cause the variable to leave that range. A

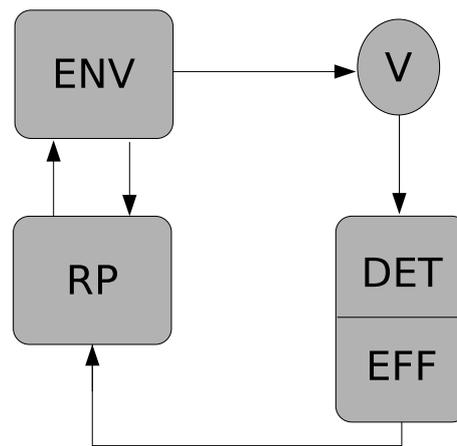


Figure 1.4: A homeostatic system consists of a reacting part (*RP*) which is continually interacting with its environment (*ENV*). This interaction determines the values of a set of essential variables (*V*). The regulating part consists of a detector mechanism (*DET*) which detects deviations of the essential variables from their target range, and an effector mechanism (*EFF*) which acts on the parameters of the reacting part. If the regulator is effective the essential variables are brought back into bounds after any deviation. There are two feedback loops, one the continual interaction between the reacting part and its environment, and the other an intermittent feedback active when the essential variables are out of bounds.

homeostatic system consists of a reacting part by which the system interacts with its environment, a set of essential variables which are the targets of regulation, a target range for each essential variable that defines the boundaries of homeostasis, and a set of regulatory mechanisms which maintain the essential variables within bounds. Every homeostatic system is co-defined with a source of perturbations that would cause a loss of homeostasis of essential variables if not counter-acted by the regulatory mechanisms. This scheme is shown in Figure 1.4, and is related to Ashby's concept of the ultrastable system (see Chapter 2).

Homeostasis is achieved when all essential variables are inside the target range, and homeostasis is lost when an essential variable moves outside of the target range. A definition of homeostasis as the static condition of having all essential variables inside the target range would be forced to label a system where essential variables went temporarily out of bounds as non-homeostatic. By emphasising the dynamic nature of homeostasis and viewing it as a process, we avoid such problems. Our definition of a homeostatic system is tolerant of the situation where an essential variable goes temporarily outside of the target range but is then returned to the target range by the action of the regulatory mechanism, so that a system where this occurs is still defined as a homeostatic system.

Homeostasis involves two feedback loops. One is the feedback loop engendered by

the continuous interaction between the reacting part of the system and its environment, which is always active because the system is situated in the world. The other feedback loop is only operational when the essential variables go out of bounds, and involves the regulator detecting this deviation and triggering a response in the reacting part which may bring the essential variables back into bounds.

Each regulatory mechanism can be sub-divided into detector and effector sub-mechanisms. Detectors discriminate between homeostatic and non-homeostatic states of the system, while effectors change the behaviour of the reacting part so that homeostasis is restored.

The environment is the source of perturbations to the system and is also a 'black box' which translates the actions of the reacting part into the values of the essential variables. It should be noted that the environment of a homeostatic system is not necessarily the same as the environment of any larger system of which the homeostatic system is a part. For example, the environment of the thermoregulatory system in mammals includes the external environment as well as most of the body, i.e., all system components that are not directly active in temperature regulation. The interaction between the active components which form the reacting part of the homeostatic system (such as sweat glands, fat-burning intracellular reactions, etc.) and the environment is what determines the temperature of the animal. The regulator mechanism acts through the reacting part, which in turn acts through the environment to determine the essential variables.

Homeorhesis is a special case of homeostasis where the target ranges for essential variables may change over time. Regulation in a homeorhetic system occurs around a trajectory rather than a set point. For example, the automatic pilot in an aeroplane corrects any deviation from the plotted course between two airports. This course will involve sections with different compass headings and angles of ascent, so that the target of regulation is constantly shifting.

1.3.4.2 Homeostatic adaptation

Homeostatic adaptation is the process by which a system changes its behaviour so that homeostasis is recovered after perturbation. Such adaptation may involve temporary responses in a system with no persistent effect, or may involve changes with a lasting effect on system parameters. In order to differentiate the two classes of adaptation, we here define *elastic homeostatic adaptation* as adaptation that does not cause any persistent changes in the system, and *plastic homeostatic adaptation* as adaptation where there is a persistent change in at least one part of the system. Elastic homeostatic adaptation means that a system will return to its original state after a perturbation is removed, whereas plastic homeostatic adaptation implies that the system will remain in an altered state after the

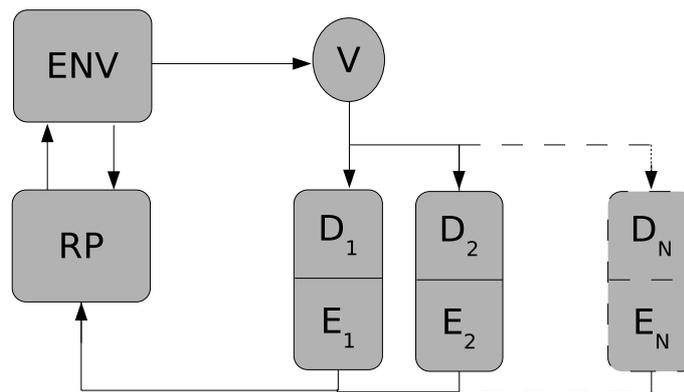


Figure 1.5: A homeostatic system can have multiple homeostatic subsystems operating concurrently. These may operate on the same essential variables as other homeostatic mechanisms, but need not have the same target ranges. This allows for different mechanisms to be called into operation for deviations of different extremity. All homeostatic feedback loops act on the same reacting part and environment. The diagram shows multiple different regulator mechanisms ($D_1E_1, D_2E_2, \dots, D_NE_N$) creating multiple homeostatic feedback loops through the essential variables (V), the reacting part (RP), and the environment (ENV).

perturbation is removed.

We make a distinction between these two types of homeostatic adaptation to avoid any confusion that might otherwise arise from discussion of ‘adaptation’, ‘adapting’, ‘adapted’ and ‘adaptive’ behaviours. We define homeostatic adaptation as any process by which a system changes its behaviour so that it moves from a non-homeostatic state to a homeostatic state, with plastic and elastic forms of homeostatic adaptation defined as above.

1.3.4.3 Multiple homeostatic feedbacks

Multiple homeostatic feedback loops can regulate the same set of essential variables in a homeostatic system. These may have (but do not require) different target ranges for the variables, and are differentiated by their regulatory mechanisms (i.e., their detector-effector mechanisms). All these homeostatic mechanisms act on the reacting part, except in the special case of higher order regulation (see next section). This scheme is shown in Figure 1.5.

1.3.4.4 Second order homeostasis

A homeostatic regulatory mechanism may itself be the subject of a higher order homeostatic feedback loop. In this situation a regulatory mechanism acts on the parameters of

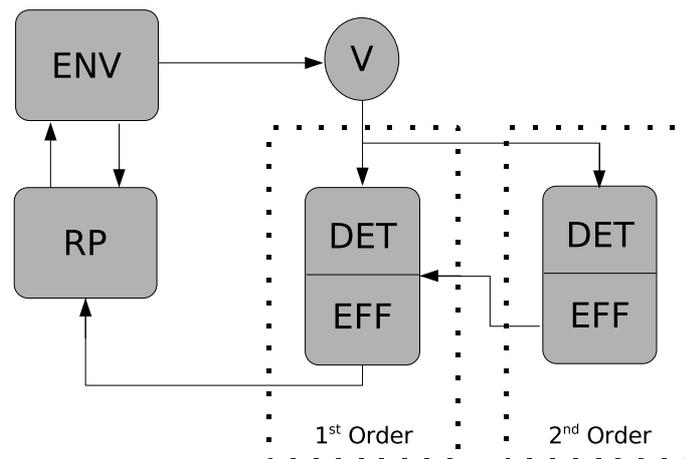


Figure 1.6: Higher order homeostatic feedbacks are possible in which the target of the higher order homeostatic mechanism is itself a homeostatic mechanism at a lower level. In this situation, all homeostatic mechanisms still monitor the same set of essential variables, but the effector output of the higher order regulator mechanism acts on a lower level regulator. This kind of arrangement can trigger changes in how regulation is accomplished when a major failure of homeostatic regulation is detected. The diagram here shows a second order regulator mechanism acting on a first order regulator, creating two tiers of homeostatic feedback loops.

another regulatory mechanism. First order homeostatic mechanisms act on the reacting part, while second order mechanisms act on first order mechanisms. This definition can be extended to the N^{th} order. Both first and second order mechanisms are triggered by changes in the same essential variables. A higher order mechanism is called into play when there is a failure of regulation by a lower order mechanism, which may be signified by a continuing loss of homeostasis or a deviation of an unusually large magnitude. This scheme is shown in Figure 1.6.

1.3.5 Examples from a minimal homeostatic system

Consider a simple system where the reacting part has a single state variable x and the environment is experienced as another variable y which changes periodically. The essential variable results from the interaction between reacting part and environment and is defined as the product xy . The regulator is composed of a number of conditional rules on the value of the essential variable. The 'IF' part of the rule corresponds to the detector in Figure 1.4, while the 'THEN' clause corresponds to the effector. These rules alter the value of x dependent on the value of xy . This system is shown in Figure 1.7.

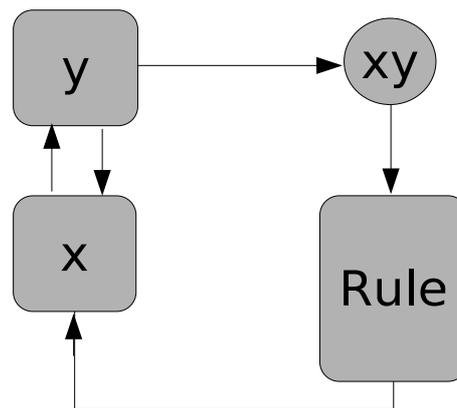


Figure 1.7: A minimal homeostatic system. The reacting part has a single state variable x which interacts with the environment, represented by a random variable y , to form a value xy for the essential variable. A conditional rule changes x dependent on xy . Depending on the form of this rule, the system can display various forms of homeostatic or homeorhetic adaptation.

1.3.5.1 An elastic homeostatic adaptive system

Let $x = 5$ at initialisation and let y be a randomly drawn integer value from the range $[0, 10]$, where y is periodically randomly reassigned. Then if we define the rule:

R: If $xy < 40$ then $x = x + 1$
 Else if $xy > 60$ then $x = x - 1$
 Else $x = 5$

we have an elastic homeostatic adaptive system where the target range for the essential variable xy is $[40, 60]$. Outside this range directed changes are made to x such that xy is adjusted in the correct direction to return to the target range. Since x is reset to 5 whenever the system is returned to its homeostatic bounds, there are no lasting changes to the system resulting from a perturbation and the system displays elastic homeostatic adaptation.

1.3.5.2 A plastic homeostatic adaptive system

Let $x = 5$ at initialisation and let y be a randomly drawn integer value from the range $[0, 10]$, where y is periodically randomly reassigned. Then if we define a variation on the previous rule:

R: If $xy < 40$ then $x = x + 1$
 Else if $xy > 60$ then $x = x - 1$

we have a plastic homeostatic adaptive system with the same target range as the previous system. In this case the value for x is *not* reset to 5 when the essential variable returns to bounds, so any changes in x are persistent and the system shows plastic homeostatic adaptation.

1.3.5.3 A plastic homeorhetic adaptive system

Let $x = 5$ at initialisation and let y be a randomly drawn integer value from the range $[0, 10]$, where y is periodically randomly reassigned. Further, let z be a variable governed by the rule $z_t = t$ where t is the time from system initialisation. Then if we define a rule:

R: If $xy < (z - 10)$ then $x = x + 1$
 Else if $xy > (z + 10)$ then $x = x - 1$

we have a plastic homeorhetic adaptive system. Changes made by the adaptive mechanism are persistent, so the system shows plastic adaptation. The target range for the essential variable xy is $[z - 10, z + 10]$ for a constantly increasing z . Since the target for regulation changes over time, the system is homeorhetic, rather than homeostatic.

1.3.5.4 A homeostatic adaptive system with multiple homeostatic feedback loops

Let $x = 5$ at initialisation and let y be a randomly drawn integer value from the extended range $[0, 100]$, where y is periodically randomly reassigned. Then the rule set:

R1: If $xy < 200$ then $x = x + 1$
 Else if $xy > 250$ then $x = x - 1$

R2: If $xy < 50$ then $x = x + 5$
 Else if $xy > 400$ then $x = x - 5$

creates a plastic homeostatic adaptive system with multiple homeostatic feedback loops. $R1$ defines one homeostatic loop with a target range for xy of $[200, 250]$, while $R2$ defines a second loop with a more relaxed target range for xy of $[50, 400]$. This system would regulate xy homeostatically to the inner range of $[200, 250]$, since $R1$ will continue to be active until this range is reached. The $R2$ loop will only be active when the essential variable xy is at very high or very low levels, but has a more powerful adaptive mechanism which will bring xy inside its target range more quickly. Both the $R1$ and $R2$ homeostatic loops show plastic homeostatic adaptation since changes they make to x are persistent.

1.3.5.5 A second order homeostatic adaptive system

Let $x = 5$ at initialisation and let y be a randomly drawn integer value from the range $[0, 100]$, where y is periodically randomly reassigned. Then the rule set:

R1: If $xy < 200$ then $x = x + \delta$

Else if $xy > 250$ then $x = x - \delta$

R2: If $xy < 50$ or $xy > 400$ then $\delta = 5$

Else $\delta = 1$

creates a second order homeostatic adaptive system. $R1$ specifies a first order homeostatic mechanism that will adjust x to move xy towards the target range of $[200, 250]$. If $R1$ is ineffective and xy goes outside of the range $[50, 400]$, $R2$ will amplify the adjustment to x by a factor of 5 in order to return xy to the target range more quickly. $R1$ is a first order mechanism since the rule operates directly on the variable x . $R2$ is second order since it operates on the parameters to $R1$. Also, $R1$ is a plastic homeostatic adaptation process since the changes it makes to x are persistent, while $R2$ is an elastic homeostatic adaptation process since it makes a temporary change to the parameters of $R1$.

1.4 Thesis outline

1.4.1 Research aims

The aim of this thesis is to study three example homeostatic adaptive systems: Ashby's Homeostat, neural homeostatic plasticity, and the biotic regulation of the environment. These studies will be at two levels. Primarily the focus will be on exploring each of the systems in its own right, seeking to gain an understanding of its operation and further scientific knowledge of its behaviour. A secondary focus will be on comparative analysis of all of the systems, seeking to identify their similarities and differences, and to gain a more general understanding of homeostatic adaptation as a cybernetic phenomenon.

The parts of the thesis will therefore have a dual existence. On the one hand, the sections looking at each of the three example systems are designed to be self-contained units, with their own research questions, experiments and conclusions. On the other hand, each section will also be of relevance to the encompassing whole. The work presented in each section will be guided by the material. Rather than force the research in a section to fit the straightjacket of the greater aim, it is the intention here to follow the loose threads of each topic to their logical conclusion. However, it is also the intention at the end of

the thesis to try and relate each topic to the broader question, and to elucidate conclusions that are more generally relevant.

1.4.2 Case studies

1.4.2.1 Ashby's Homeostat

The Homeostat was an electrical device designed by Ashby to demonstrate the principle of ultrastability [7]. It consisted of four magnetic units connected by circuits with variable properties. Each unit was attached to a gauge, and each gauge had a target region for its level of activation. When the activation of a unit was outside of this range the parameters of the connecting circuit to that unit were altered at random, thus changing the way in which it was influenced by output from the other three units. Left to itself the Homeostat would eventually settle to a stable equilibrium where the output of all units was in bounds, i.e., where homeostasis was maintained. The outputs of the units were the essential variables, and the system as a whole was ultrastable. The Homeostat will be explored in more detail in Part II.

We develop a simulated Homeostat that demonstrates how a homeostatic adaptive system may be ultrastable, and how stability is affected by the size of the system and the tightness of the constraints. The ability of an ultrastable system to accommodate perturbations by homeostatic adaptation is also explored.

1.4.2.2 Neural homeostasis

It is increasingly recognised by neuroscientists that plastic mechanisms act in the brain to homeostatically regulate levels of neural activity [37, 167]. These mechanisms alter the properties of neurons and neuronal networks, such as synaptic connection weights and the intrinsic excitability of neurons, so that while there may be short term fluctuations in neural activity, the long term mean firing rate of each neuron tends towards some set point.

Homeostatic plasticity is a relatively new area of study in the neurosciences. It is thought to play an important role in counter-acting the destabilising positive feedback effects associated with other forms of neural plasticity such as Hebbian learning, but its precise functioning and mechanisms are yet to be fully elucidated. The network-level effects in particular remain to be explored.

Here we approach homeostatic plasticity from the viewpoint of neural robotics, which allows the study of its effects at a variety of different levels, from individual nodes through to a complete behaving agent interacting with the world. Using a highly abstracted neural

model, we demonstrate that homeostatic plasticity makes nodes and networks more sensitive to input, improves signal propagation, and makes oscillatory dynamics more likely. From an engineering standpoint, we show that it may have both beneficial and detrimental effects on the evolvability of networks for robot control. While the simplicity of the neural model employed here is well removed from its biological counterpart, we draw on our results to speculate on the functioning and role of homeostatic plasticity in biological neuronal networks, suggesting that the current neuroscientific view may be incomplete.

1.4.2.3 Environmental regulation by the biota

The Gaia hypothesis [116], which suggests that the biota collectively regulate the environment to conditions suitable for life, has been the cause of much controversy. Much of the debate concerns the compatibility of Gaia theory with evolutionary biology. The idea that the biota can evolve to collectively regulate the environment seems to run counter to selfish Darwinian adaptation, implying teleology or group selection. Some of the criticisms of Gaia have been rebutted, while others remain. The central question is how global regulation might be explained by selection acting on individuals, and it is a question that has yet to be answered.

With the aim of studying environmental regulation by the biota in general, we here approach this topic by looking primarily at Gaia theory. Specifically, we review existing models and then develop a simplified version of the Daisyworld [183] model. We use this model to replicate known Daisyworld results and to generate some new results concerning the importance of evolutionary constraints for regulation. The effects of relaxing the core assumptions of the Daisyworld model are explored, showing that successful environmental regulation relies on several key assumptions which may or may not be present in nature. A new model, the Flask model, is then developed with the intention of demonstrating the compatibility of environmental regulation with Darwinian evolution in a model which relaxes several of the key Daisyworld assumptions. Preliminary results from the Flask model are given, allowing a strong argument to be constructed for the evolution of regulatory feedbacks.

1.4.3 Methodology

The experimental work presented in this thesis relies on simulation modelling and some justification of this approach is given here.

Computer simulations are no longer a new technique and have been usefully employed in many situations. We can categorise simulation models as falling into two main classes,

which we can call ‘engineering simulations’ and ‘computational thought experiments’. It is the latter class which concerns us here.

Engineering simulations are the least controversial of the two types of simulation model. They are typically used to optimise some feature of a real world system *in silico* before its physical construction. For example, the aerodynamics of a car can be tested in simulation prior to more costly wind tunnel experiments. Engineering simulations use accepted knowledge about the world to determine the implications of a particular configuration of a system. They derive new data, without necessarily generating new knowledge.

Computational thought experiments are a more recent use of computer models and are often regarded with more suspicion than engineering simulations. The basic idea is to test the logical validity of a particular set of assumptions about how a real world phenomenon occurs. For example, a researcher might think they understand how flocking behaviour occurs in birds and decide to test their model by instantiating it as a simulation. If the simulation is based on sound assumptions and flocking behaviour is observed, then the researcher’s theory of flocking is logically consistent and becomes a live hypothesis for testing by a field biologist. If flocking behaviour is not observed in the simulation, the assumptions of the model are not sufficient and the model must be revised.

The use of thought experiments using pen and paper is an accepted part of scientific progress. Thought experiments using computer simulations are no more and no less valid than their armchair equivalents. The utility of simulation models lies in their ability to capture more complex phenomena than is possible with pen and paper. The computational power of a simulation allows more entities to be modelled and easier testing of different sets of parameters. Simulation models have been described as ‘opaque thought experiments’ [51], recognising that the phenomena they are often used to model are too complicated to be tractable by conventional thought experiments.

It should be remembered that simulation models can never prove any theories about the real world, but also that this is not their aim. Their utility is in forcing the explicit statement of the assumptions of a theory and in allowing the easy observation of the implications of these assumptions. The logical consistency of a theory can be tested and refined by simulation modelling, in order to create valid hypotheses for empirical testing.

The complexity of the systems studied in this thesis necessitates the use of simulation models. The Homeostat, neural networks, and the Earth system, are all complex systems that are opaque to armchair thought experiments.

In the case of the Homeostat, the model Homeostat we develop is not a true simulation, since it does not pretend to model some real world system but is an actual instantiation of a simple ultrastable system. The same partially applies to the neural networks used in Part III. While inspired by biology, they are in themselves real instances of a class

of computer program often used for robot control, at the same time as being simulacra of biological neuronal networks. The Gaia models used in Part IV are more traditional simulations in that their interest comes from their interpretation as simulacra of real biological systems.

1.4.4 Thesis plan

This thesis is divided into five parts which reflect the logical structure of the material presented. Parts II, III and IV cover the three case studies and are intended to be self-contained. Part II covers the Homeostat, Part III covers neural homeostatic plasticity, and Part IV covers environmental regulation by the biota. Each case study begins with a review of the relevant literature and background material, before outlining the research questions to be answered. This is followed by experimental chapters presenting the methods used and results obtained, and then some discussion of the results presented.

Part I and Part V are more concerned with the over-arching theme of the thesis: homeostatic adaptive networks. Part I consists of the current introductory chapter, which sets the scene and gives a general overview of the thesis. Part V contains discussion of the material presented in Parts II, III and IV, and analyses these results in the context of a general theory of homeostatic adaptive networks.

Part II

The Homeostat

Chapter 2

The Homeostat

2.1 Overview

The Homeostat was developed by Ashby in the 1950s as a demonstration of his concept of *ultrastability* [7], which is the earliest exposition of a theory of homeostatic adaptation. We begin by introducing the concept of ultrastability and Ashby’s view of homeostatic adaptation, including the Homeostat. Next we present a simple simulated Homeostat that is used to explore ultrastability and homeostatic adaptation. We show that the time taken for the simulated Homeostat to reach stability is affected by the number of nodes and the tightness of the homeostatic constraints on the nodes. The Homeostat can adapt to perturbations and is in some circumstances capable of being used as a self-organising control system.

2.2 Background

William Ross Ashby (1903-1972) was a British cybernetician originally trained in neurology. Ashby’s most famous work is *Design for a Brain* [7], in which he attempts to answer the question of how the brain can produce adaptive behaviour:

“When a kitten first approaches a fire its reactions are unpredictable and usually inappropriate. It may walk almost into the fire, or it may spit at it, or may dab at it with a paw, or try to sniff at it, or crouch and ‘stalk’ it. Later

however, when adult, its reactions are different. It approaches the fire and seats itself at a place where the heat is moderate. If the fire burns low, it moves nearer. If a hot coal falls out, it jumps away. Its behaviour towards the fire is now ‘adaptive’.” (p.12, [7])

Ashby looked at organisms as ‘machines’, that is, he regarded them as determinate and subject to the same physical laws as mechanical devices. He defined the *essential variables* of an animal as those variables which must be kept within limits for the animal to remain alive. Armed with this definition, Ashby could recast the problem:

“A determinate machine changes from a form that produces chaotic, unadapted behaviour to a form in which the parts are so co-ordinated that the whole is stable, acting to maintain its essential variables within certain limits - how can this happen?” (p.70, [7])

This question has obvious implications for learning in artificial agents as well as animals, and in recent years there has been a resurgence of interest in Ashby’s work [46,48].

2.2.1 Ultrastability

Ashby developed the idea that organisms are *ultrastable*. He made a functional separation between different parts of the organism that specified a set of essential variables, a reacting part, and a set of parameters to the reacting part (see Figure 2.1). The essential variables V must be kept within bounds to ensure the continued viability and survival of the organism. The reacting part R was the behaviour-producing part by which the organism interacted with the world. The parameter set S determined how R should react to the environment.

Ultrastability was then defined as the situation where:

“Two systems of continuous variables (that we call ‘environment’ and ‘reacting part’) interact, so that a primary feedback (through complex sensory and motor channels) exists between them. Another feedback, working intermittently and at a much slower order of speed, goes from the environment to certain continuous variables which in their turn affect some step-mechanisms, the effect being that the step-mechanisms change value when and only when these variables pass outside given limits. The step-mechanisms affect the reacting part; by acting as parameters to it they determine how it shall react to the environment.” (p.98, [7])

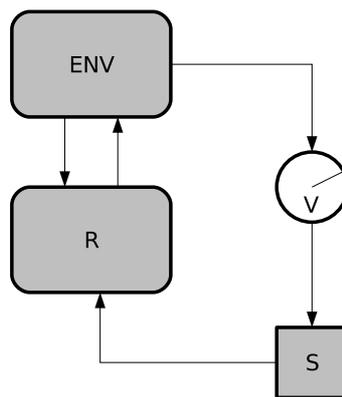


Figure 2.1: Ultrastability. There are two feedback loops: the primary loop involves the interaction of the reacting part R with its environment ENV , while a secondary loop involves also the essential variables V and the parameters S to R . When V goes out of bounds step-changes are triggered in S so that the behaviour of R changes. This affects the interaction between R and ENV , which may or may not bring V back in bounds.

The ‘step-mechanisms’ Ashby refers to are mechanisms that cause sudden discontinuous changes in the value of a variable. Thus, in other words, an ultrastable system is one where an excursion of the essential variables from their desirable ranges causes a sudden change in the parameters of the behaviour-producing subsystem. An ultrastable system can only be stable when normal operation maintains the essential variables within bounds, since if this does not happen the system is made unstable by step-changes in parameters. Stability requires homeostasis of essential variables, and the lack of homeostasis triggers changes in behaviour.

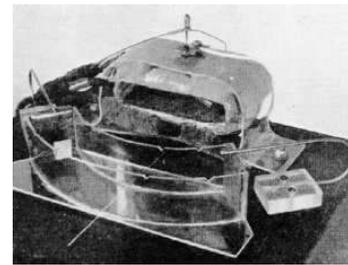
2.2.2 The Homeostat

The Homeostat was an electromechanical device that Ashby built to demonstrate ultrastability. It consisted of four units mounted on a base platform (Figure 2.2(a)). Each unit carried a pivoted magnet (Figure 2.2(b)), and the angular deviations of these magnets from the central position were specified as the essential variables of the system. Each unit emitted a DC current proportional to the deviation of its magnet, which was passed to the other three units.

The currents received by each unit acted upon its magnet via three coils, and the output of a unit also affected its magnet via a self-connection. Before reaching the coils, input currents were modified by passing them through a potentiometer and a commutator, the settings for which acted as parameters to the unit. On each unit, the potentiometer and commutator settings were assigned by a uniselector which chose from 25 different



(a) Homeostat



(b) Single unit

Figure 2.2: Ashby’s Homeostat consisted of four units mounted on a base platform. The output from each unit was fed into all other units by circuitry that was reconfigured when unit output moved outside the target range, such that the whole system was ultrastable. Illustrations adapted from [7].

random settings. This gave $25^4 = 390625$ different combinations of parameters for the four unit system. If the angular deviation of a magnet passed outside the range $[-45^\circ, 45^\circ]$ the uniselector on that unit would choose a new setting at random.

The Homeostat was shown to be ultrastable. It was initialised with the uniselectors set to random positions and its subsequent behaviour was observed. By changing the uniselector setting when magnet deviations went out of bounds, eventually a stable equilibrium was reached where all magnets were located in the target range for angular deviation. The Homeostat was also shown to be able to adapt to perturbations. If a magnet was manually displaced, the ultrastable nature of the system ensured that the system eventually settled down to an attractor where all free-moving magnets were in bounds.

2.3 Method

In this section we will develop a simple simulated Homeostat that displays the property of ultrastability. The Homeostat described here is constructed differently to the original Homeostat. Ashby’s Homeostat was an electromechanical device, whereas the current implementation is a computer program, and the connectivity and mechanisms are necessarily different in nature. However, both Ashby’s Homeostat and the current version are designed with similar logical structure and display similar behaviours, so for that reason the ‘Homeostat’ nomenclature is retained.

Our Homeostat is a system of N units arranged in a fully connected network topology. Each unit receives N inputs, from itself and from all other units, weighted by the strength of the connection between them (see Figure 2.3). The weighted sum I (Equation 2.1) of the inputs to a unit determines its output s , as specified by a piecewise linear transfer

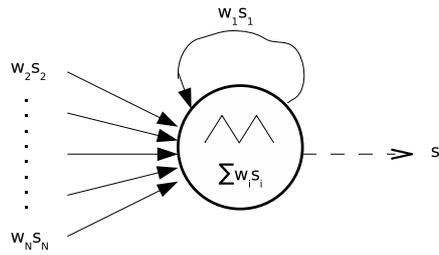


Figure 2.3: Schematic of a single Homeostat unit. The unit receives input from other units ($w_i s_i$) and from itself ($w_1 s_1$). Output s_1 is a piecewise linear function of weighted summed input.

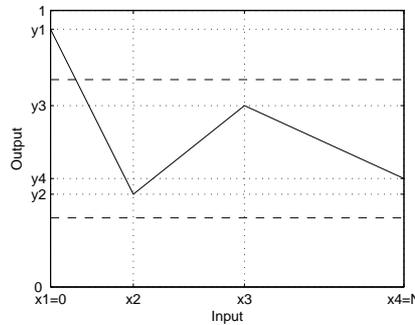


Figure 2.4: Example Homeostat unit transfer function. Function is piecewise linear with points at $(x_1, y_1), \dots, (x_P, y_P)$ where $x_1 = 0$ and $x_P = N$ for all units in an N -unit network. $P = 4$ in example shown. Dashed lines indicate target homeostatic range.

function F (Equation 2.2 and Figure 2.4). A unit may be specified as a set of parameters $U = \{w_1, \dots, w_N, x_1, \dots, x_P, y_1, \dots, y_P\}$ where w_i is the afferent connection strength from the i^{th} unit and (x_j, y_j) is the coordinate of the j^{th} point on its transfer function. Ranges are set so that $w \in [0.00, 1.00]$ and $s \in [0.00, 1.00]$, so $I \in [0.00, N]$.

$$I = \sum_i^N w_i s_i \quad (2.1)$$

$$s = F(I) = \begin{cases} y_1 + (y_2 - y_1) \left(\frac{I - x_1}{x_2 - x_1} \right) & : & x_1 \leq I < x_2 \\ y_2 + (y_3 - y_2) \left(\frac{I - x_2}{x_3 - x_2} \right) & : & x_2 \leq I < x_3 \\ y_3 + (y_4 - y_3) \left(\frac{I - x_3}{x_4 - x_3} \right) & : & x_3 \leq I \leq x_4 \end{cases} \quad (2.2)$$

At initialisation, connection strengths are randomly assigned from a uniform distribution on the appropriate range, as are values for all transfer function parameters. A target

range $R = [0.5 - \delta, 0.5 + \delta]$ for output is specified, where the size of δ determines the tightness of the homeostatic constraint. If $s \in R$ the unit is homeostatic. If $s \notin R$ then homeostasis is lost and adaptive change is triggered.

There are two adaptive mechanisms which are applied to the parameters of non-homeostatic units. The first mechanism assigns new random values for the strengths of all afferent connections to the unit (Equation 2.3). The second mechanism assigns new random values for the coordinate parameters of the unit's transfer function (Equation 2.4). Ranges for reassigned parameters are the same as those used for initialisation.

$$IF (s \notin R) THEN [w = rand(0.00, 1.00) \quad \forall w \in \{w_1, \dots, w_N\}] \quad (2.3)$$

$$IF (s \notin R) THEN [x = rand(0.00, N) \quad \forall x \in \{x_2, \dots, x_{P-1}\}] \\ AND [y = rand(0.00, 1.00) \quad \forall y \in \{y_1, \dots, y_P\}] \quad (2.4)$$

where $rand(a, b)$ is a function that returns a real number randomly drawn from a uniform distribution on the range $[a, b]$.

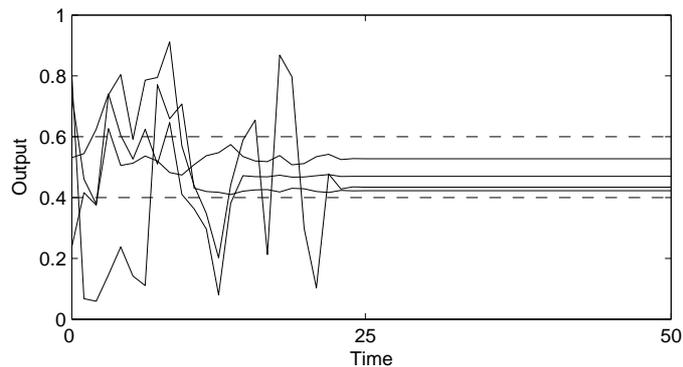
2.4 Results

Here we describe the behaviour of the simulated Homeostat. We look first of all at its ultrastable behaviour when each of the two different adaptive mechanisms is used, before looking at how the time taken for the system to reach stability is affected by the number of nodes in the network and by the tightness of the homeostatic constraint. Next we look at the response of the system to perturbation, before finally considering the possibility of using the simulated Homeostat as a self-organising control system.

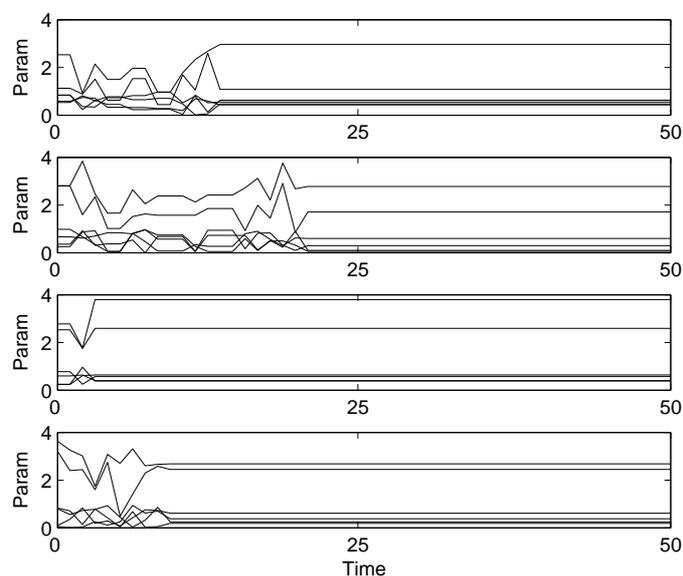
Two different adaptive mechanisms were used: a mechanism which changes transfer function parameters, and a mechanism which changes connection strengths. The two mechanisms are used to show that homeostatic adaptation can occur via the action of different kinds of mechanism, and that some mechanisms are more effective at maintaining homeostasis in the system than others. The two mechanisms can lead to different kinds of homeostatic 'solution'.

2.4.1 Reaching stability

The first result to show is that the Homeostat described in the previous section is ultrastable. Figure 2.5 shows the Homeostat converging to a stable steady state from a random initialisation. Here the target range was set to be $[0.4, 0.6]$ and a 4 unit network



(a) Unit output



(b) Transfer function parameter change

Figure 2.5: Convergence over time to a stable steady state in a 4-unit Homeostat with transfer function adaptation. Comparison of the two plots shows that parameter change is correlated with excursions of unit output from the target range. The upper plot shows the target range for unit activation ($[0.4, 0.6]$, region within dashed lines) together with the activation of all of the units (solid lines). Initially activity is out of bounds and changes erratically, but over time the activity of all nodes stabilises at a steady level within the target range. The lower plots show the change in the (x, y) coordinate parameters of the transfer function for each unit. These change continually so long as the activity of the associated unit is out of bounds, but stop changing when activity is within bounds. The lower plots therefore show a large amount of parameter change early on when activation of several units is out of bounds. Eventually the random changes to transfer function parameters generate a parameter set that keeps activity in bounds, at which point parameter change ceases and activation stabilises.

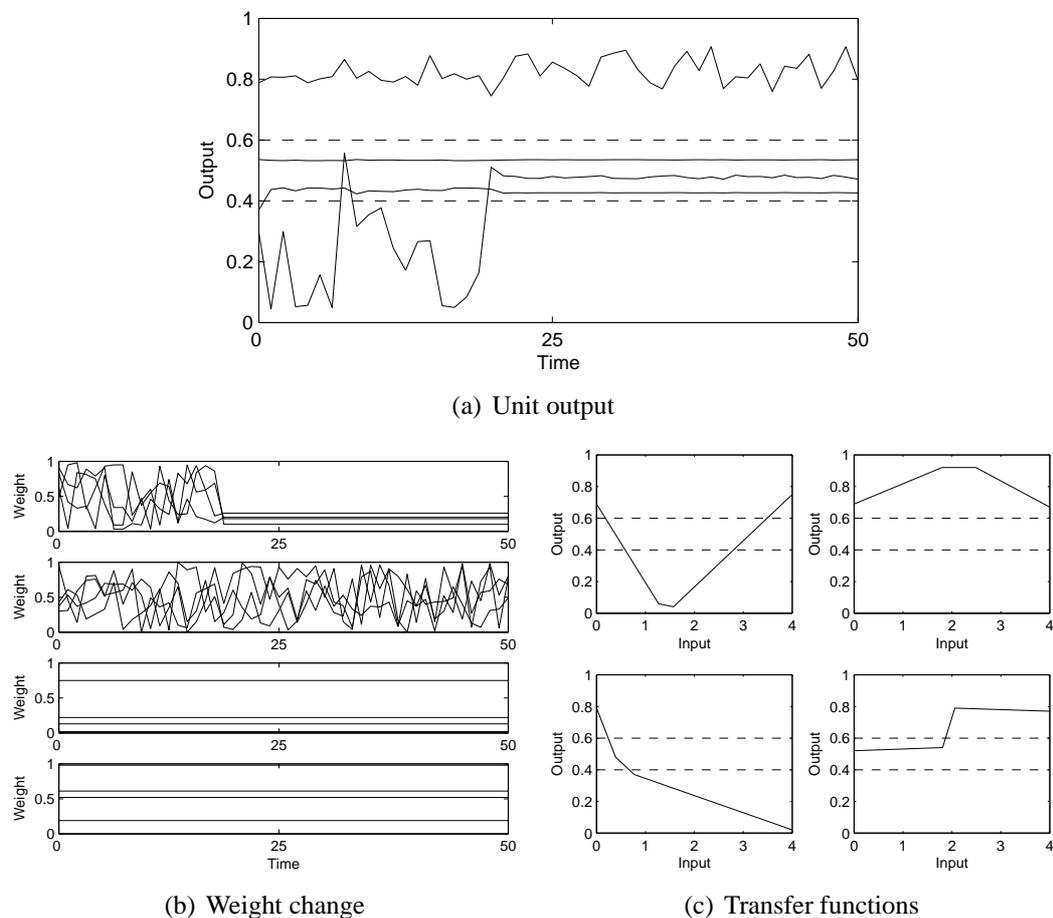


Figure 2.6: Convergence/non-convergence to homeostatic stability of units in a 4-unit Homeostat with adaptation by reassignment of afferent connection strengths. The upper plot shows activation of all units, while the lower plots show connection weights to each unit (left plot) and the fixed transfer function for each unit (right plot). Transfer functions are fixed for the duration of the trial. Comparing the plots, we can see that two of the units immediately find a good set of connection strengths that keeps their activity in bounds, so that these units show a fixed level of activity and no connection strength change (lower two weight change plots, bottom left and bottom right transfer function plots). Another unit initially has activity out of bounds but eventually finds a good set of connection strengths, at which point parameter change ceases (uppermost weight change plot, top left transfer function plot). The remaining unit never finds a good parameter set. Its transfer function (top right transfer function plot) has no part of its input range giving output in the target homeostatic range, so change in its afferent connection strengths continues indefinitely (second-from-top weight change plot). This is correlated with continually changing activity that is always outside the target range (upper line in activity plot).

was used. Output from several units starts off outside of the target range but quickly moves to the target range as new transfer function parameter sets are generated that keep activity in bounds. Parameter change can be seen in Figure 2.5(b), which shows the changes in transfer function parameters for each unit. Parameter change is correlated with excursions of unit output from the homeostatic target range.

Figure 2.6 shows the behaviour of a Homeostat where the adaptive mechanism is random reassignment of the strengths of the afferent connections to any non-homeostatic unit. In this scenario, most units find a set of connection weights that satisfy the homeostatic constraint, but some units are not able to do so because their transfer functions allow no possibility of giving output in the correct range, no matter what level of input is received. Figure 2.6(c) shows the transfer functions for each of the units in the Homeostat. The convergent units all have transfer functions where it is possible for the homeostatic constraint to be satisfied, while the non-convergent unit (top right transfer function plot) has a transfer function which precludes convergence to a stable steady state. Its output is always out of bounds.

Homeostats that adapt by reconfiguring their transfer function are always able to eventually find a set of transfer function parameters that leads to steady output in the homeostatic range. Homeostats that adapt by assigning new afferent connection strengths are not always able to satisfy the constraint, since their fixed transfer function may not allow this possibility.

2.4.2 Time to convergence

The time taken for a Homeostat to converge to a stable homeostatic attractor with the activation of all units inside the target range may be affected by the number of units in the Homeostat network and the tightness of the homeostatic constraint on each unit. In this section we examine these relationships by varying these parameters and measuring the time taken for the Homeostat to reach stability.

As a proxy for the time taken to reach a stable attractor, we measure the probability of individual Homeostat units having their activation in bounds after a given amount of time has elapsed from initialisation. This method is used because some Homeostats never converge to a homeostatic stable state (e.g., if they use connection strength adaptation and have an unsuitable transfer function for one of the units) and also because random changes to parameters can be an inefficient search method that occasionally takes a very long time to find a solution. Both of these occurrences make accurate measurement of the time taken to reach a homeostatic attractor difficult, but the measurement of probability that a unit will be homeostatic after a given time allows us to gather the information we need to look

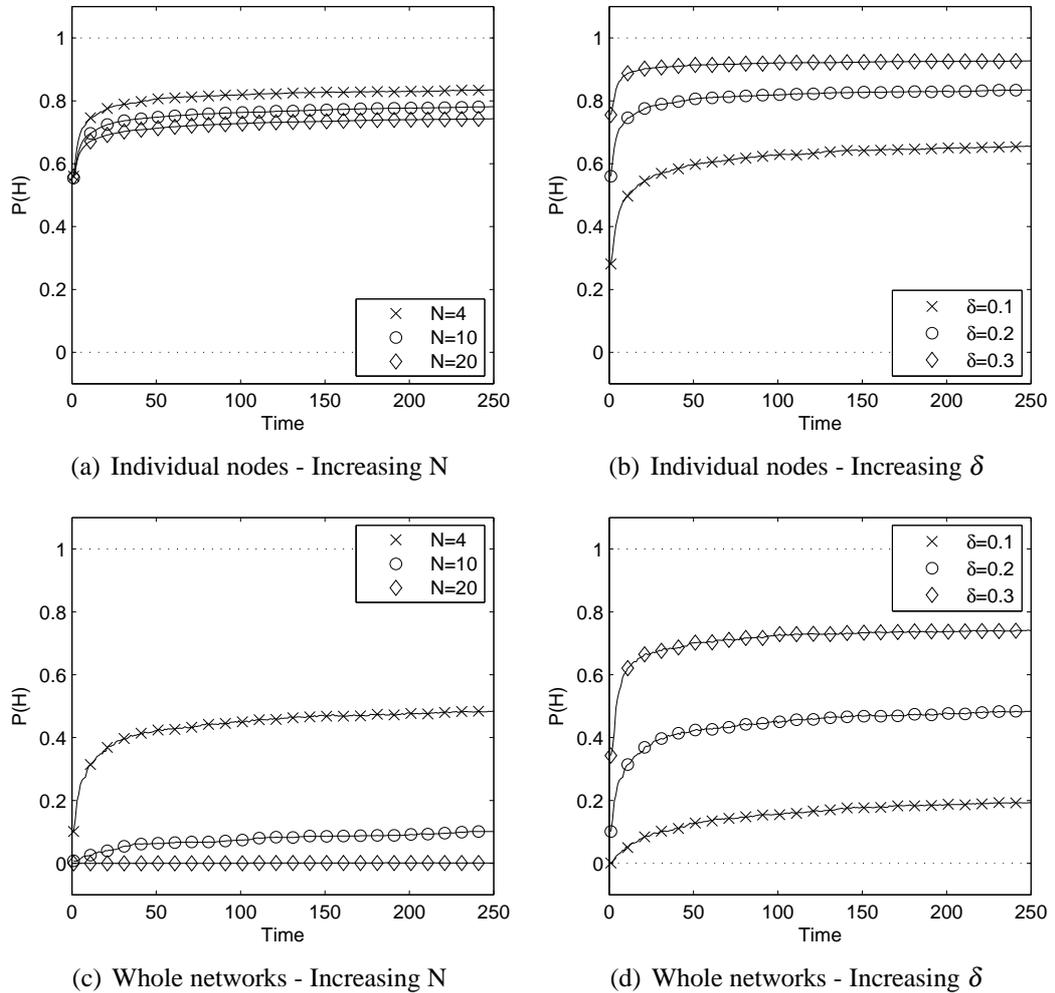


Figure 2.7: Probability of homeostasis in Homeostats with adapting connection strengths, plotted against elapsed time from initialisation. *Top row:* Individual nodes. *Bottom row:* Whole networks. Plots show the effect of increasing network size ($N \in \{4, 10, 20\}$) when homeostatic target range is kept fixed at $[0.4, 0.6]$ (left column) and the effect of keeping network size fixed at $N = 4$ and varying the tightness of the homeostatic target range ($[0.5 - \delta, 0.5 + \delta]$ for $\delta \in \{0.1, 0.2, 0.3\}$) (right column). Increasing N reduces the likelihood of nodes and networks displaying homeostasis. Relaxing the homeostatic constraint by increasing δ has a more dramatic effect, making homeostasis much more likely. All the Homeostats are unlikely to display perfect homeostasis because of the possibility of fixed transfer functions which do not allow the constraint to be satisfied, as occurred in the non-convergent Homeostat shown in Figure 2.6.

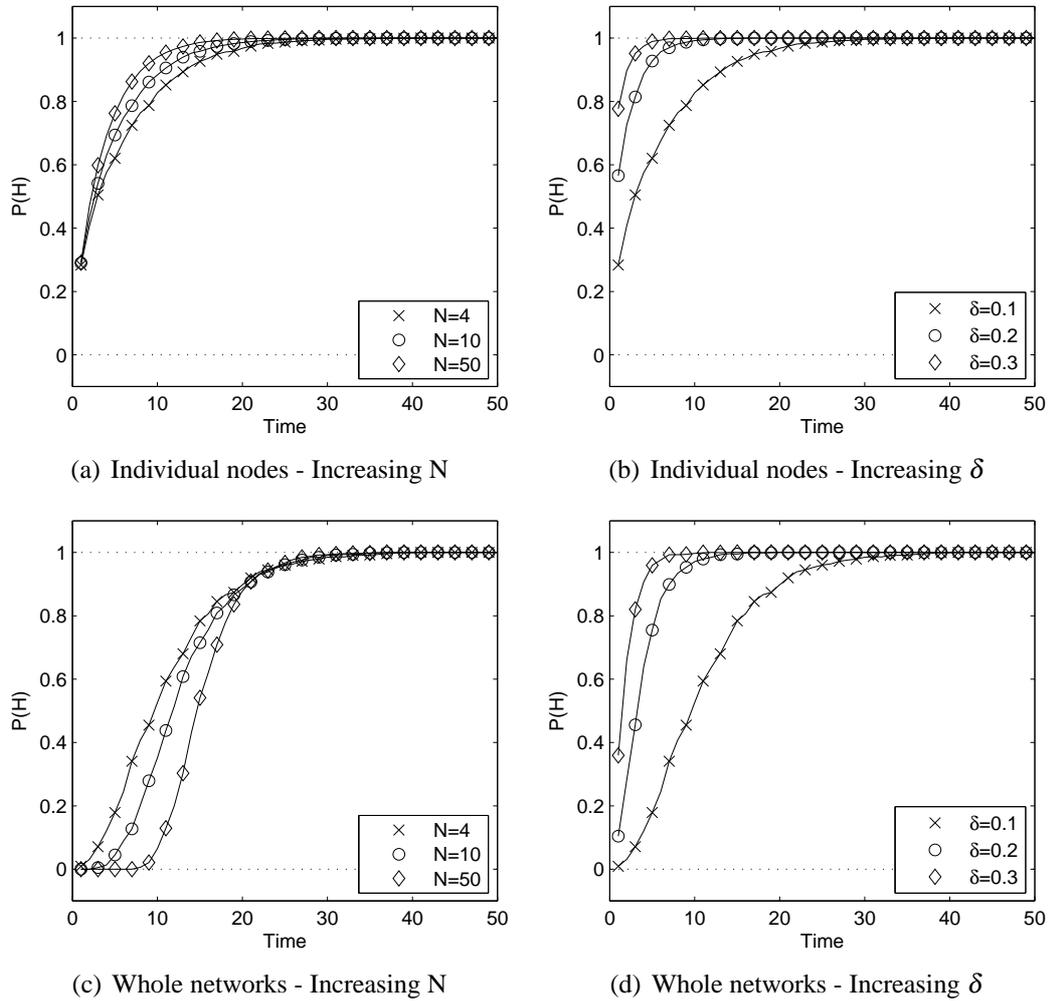


Figure 2.8: Probability of homeostasis in Homeostats with adapting transfer functions, plotted against elapsed time from initialisation. *Top row:* Individual nodes. *Bottom row:* Whole networks. Plots show the effect of increasing network size ($N \in \{4, 10, 50\}$) when homeostatic target range is kept fixed at $[0.4, 0.6]$ (left column) and the effect of keeping network size fixed at $N = 4$ and varying the tightness of the homeostatic target range ($[0.5 - \delta, 0.5 + \delta]$ for $\delta \in \{0.1, 0.2, 0.3\}$) (right column). In this scenario, increasing N counter-intuitively causes a small increase in the likelihood of individual nodes displaying homeostasis, though it also decreases the likelihood of whole networks displaying homeostasis. This is because higher connectivity reduces the relative impact of individual afferent signals to a unit by extending the size of the range of inputs for which its activity falls in the target range. As δ is increased and the homeostatic constraint is relaxed, the probability of individual nodes and whole networks displaying homeostasis goes up, as would be expected. Perfect homeostasis is observed after around 30 timesteps, because it is always possible to find a transfer function that can satisfy the homeostatic constraint by varying the (x, y) coordinates.

at the effect of different Homeostat parameters on speed of convergence. Because we are only interested in the qualitative changes resulting from different parameterisations, the method used here is sufficient for our purpose, despite its poor handling of cases where convergence does not occur.

We look at the performance of both the adaptive transfer function mechanism and the adaptive connection strength mechanism that were used in the previous section. Homeostats are initialised with random parameters and then updated for a fixed period in which the proportion of units displaying homeostasis is measured at each timestep. These data are used to calculate the probabilities of individual units and of entire Homeostat networks showing homeostasis after a given time. This information is gathered from 1000 Homeostats of each type examined.

Figure 2.7 shows the effect of changing the number of units or the tightness of the target range in Homeostats that adapt by random reassignment of connection strengths. Figure 2.8 shows the effect of the same changes on Homeostats that adapt by random reassignment of transfer function parameters. Comparison of Figures 2.7 and 2.8 shows that Homeostats using the adaptive transfer function mechanism are much quicker to stabilise than Homeostats using the adaptive connection weight mechanism, with both nodes and networks having a near unity chance of having converged after around 30 timesteps when the adaptive transfer function is used. The adaptive connection weight mechanism is hampered by the occurrence of transfer functions which do not permit constraint satisfaction, which has an adverse effect on the likelihood of reaching homeostasis.

Increasing network size reduces the likelihood of a whole network behaving homeostatically for both mechanisms, and has a similar effect on individual nodes when the adaptive connection strength mechanism is used. However, increased network size actually increases the likelihood of homeostasis in individual nodes when the adaptive transfer function is used. This effect is seen because the size of the ‘homeostatic’ input range increases in proportion to the number of units in the network, which reduces the effect of changes in input on unit output. Such an effect would also be seen in the case of the adaptive connection strength mechanism, but is masked by the greater impact of the problem of fixed transfer functions that are incompatible with homeostasis.

The tightness of the homeostatic constraint has a more significant, but more straightforward effect on the probability of homeostasis. As the target range for unit output gets smaller, increasing the tightness of the constraint on the Homeostat, both nodes and networks take longer to reach homeostasis. Again the adaptive transfer function mechanism is more effective than the adaptive connection weight mechanism.

The effects of network size and the tightness of the homeostatic constraint on the time and probability of convergence to a homeostatic stable state are of interest primarily for

pedagogical purposes, serving to illustrate some of the factors which may influence the ability of a system to maintain homeostasis. However, the results may have implications for systems other than the Homeostat. While caution should be exercised in seeking to transfer insight from the Homeostat to other systems, the observations that bigger networks and tighter constraints make homeostasis harder to achieve in the Homeostat seem intuitively likely to hold in many kinds of system.

2.4.3 Perturbations

The Homeostat can only really be said to be maintaining homeostasis in its constituent units if there is some form of external perturbation that threatens to push the target variables out of bounds. In this section we describe the behaviour of the Homeostat when it is subjected to a variety of different perturbations. Perturbations take the form of an external input to each Homeostat unit, implemented as an amount added to the weighted sum of the inputs it receives from other units. For these experiments, only the adaptive transfer function mechanism is used, since it has been shown to be the most effective form of adaptation.

Figure 2.9 shows the effect of perturbations applied to a single unit in the Homeostat. A random amount drawn from a uniform distribution in the range $[0, 2]$ is added to the input of a designated unit in a 4-unit Homeostat. Initially, the perturbed unit often loses homeostasis and undergoes parameter change, but over time a more stable parameter set is found that allows the unit to maintain homeostasis in the face of continuing perturbations. Sometimes the change in output from the perturbed unit disrupts other units in the network, and causes a cascade of parameter change, but again this effect is reduced over time. The perturbed unit adapts so that it nullifies the effect of the perturbation on the Homeostat.

Figure 2.10 shows the effect of simultaneous perturbations applied to all units in the Homeostat. Two forms of perturbation, one which adds a random level of external input to all units, and one which adds a continuously varying sinusoidal input signal to all units. The Homeostat adapts to both forms of perturbation, with all units eventually finding sets of transfer function parameters that allow them to maintain homeostasis in the face of continuing perturbation.

2.4.4 A control problem

We have seen that the Homeostat can maintain homeostasis and adapt to perturbations, and in this section we will now look at whether ultrastability can be used to create useful controllers. We use a simple input-output mapping task in which the Homeostat must

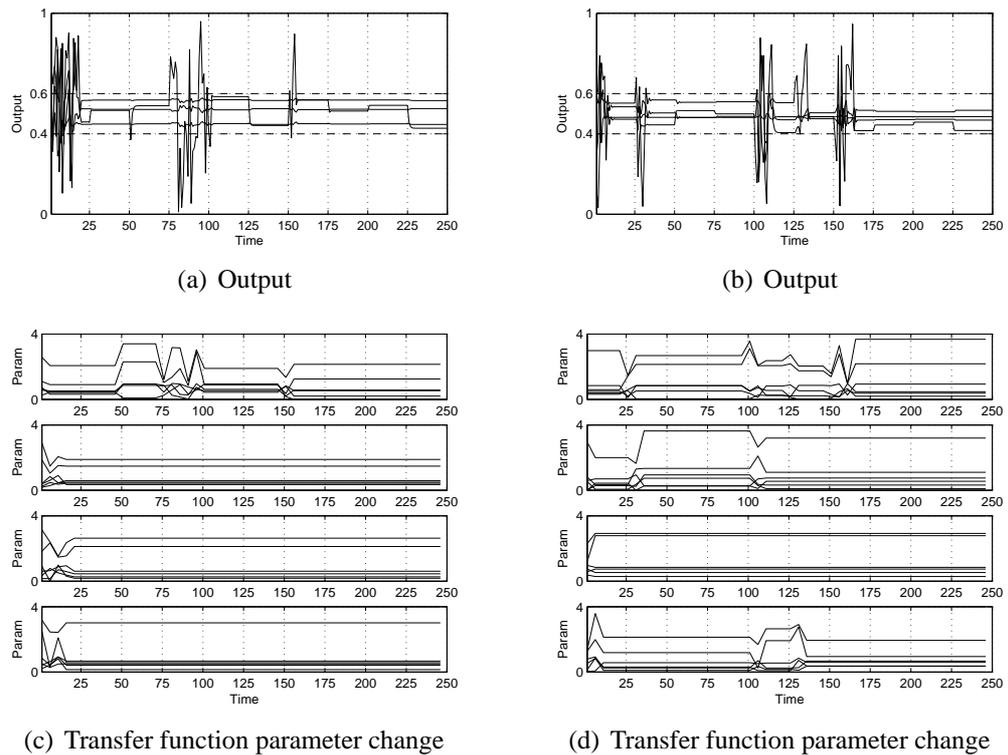


Figure 2.9: 4-unit Homeostat with adaptive transfer function where random input from the range $[0, 2]$ is applied only to a single designated input unit. Input level changes every 25 timesteps. Changes in input can cause activity of the input unit to go out of bounds, causing adaptive change. Sometimes this can lead to a cascade of changes around the network, where a perturbation causes loss of homeostasis in the input unit, which in turn causes a loss of homeostasis and adaptive change in other units. *Left column:* Case where perturbations only cause loss of homeostasis in the input unit. *Right column:* Case where perturbation causes a cascade of homeostasis-loss around network.

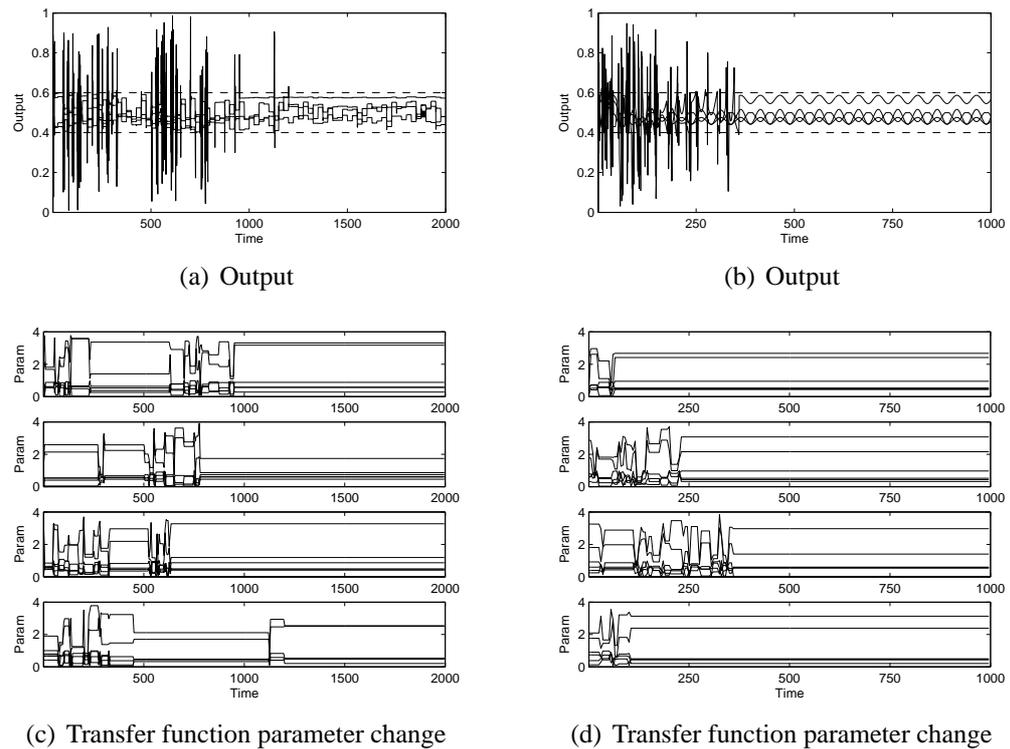


Figure 2.10: 4-unit Homeostat with adaptive transfer function with perturbation applied to all units. Examples are shown where the Homeostat adapts to two different forms of perturbation. *Left column:* Random input from the range $[0, 2]$ is applied to all units, changing every 50 timesteps. All units eventually find transfer function parameters to accommodate this perturbation and maintain homeostasis. *Right column:* Sinusoidal input is applied to all units with amplitude $[-1, 1]$ and period 50 timesteps. All units eventually adapt to this perturbation.

give the correct response to different input signals. Two levels of external input are applied to the units of the Homeostat, which must respond with the correct level of output from a designated output unit. If the output response is incorrect, change is triggered in the parameters of the output unit. All other units must maintain homeostasis as before, with parameter change applied if their activation goes out of bounds. The task for the Homeostat is thus to give the correct output associated with each input signal, while maintaining homeostasis. Again, only the adaptive transfer function mechanism is used for these experiments.

Figure 2.11 shows results when the input signal is applied to all units, including the output unit. There were two levels of input signal: high (external input to all nodes is equal to 2), and low (external input to all nodes equal to 0). The correct output response to the high input signal was output in the range $[0.7, 0.8]$, while the correct response to the low input signal was output in the range $[0.2, 0.3]$. Effectively this means that the output node had a split target range for homeostasis; when a low response was required the target range was $[0.2, 0.3]$ and when a high response was required the target range was $[0.7, 0.8]$. Should the output unit give the wrong response, i.e., a response outside whichever target range was currently being applied, adaptive change was triggered. Non-output nodes had a target homeostatic range of $[0.4, 0.6]$. The input signal was changed every 50 timesteps.

Figure 2.11 shows that the Homeostat successfully achieved the task. After a period of adaptation a parameter set is found that gives correct output responses to changes in input and maintains homeostasis in all other nodes.

However, in a variation of the experiment where the input signal was applied to all units *except* the output unit (meaning that a signal must be passed across the network) it proved extremely difficult to satisfy the control problem. This is because the homeostatic nature of the Homeostat units means that they tend not to be very responsive to input, since the units self-organise so that output always lies within a narrow target range. The activation of the input units does not change significantly enough in response to changing external input to cause a change in the activation of the output unit. Units in the Homeostat are better at stopping information spreading than propagating signals.

2.5 Discussion

In this chapter we have presented a simple Homeostat inspired by Ashby's original device, that demonstrates the concept of ultrastability. Ultrastable systems are those in which a threat to homeostasis of essential variables is countered by adaptation of the interacting part such that homeostasis is maintained during normal behaviour. The simulated Homeostat we developed here maintains homeostasis in the levels of activation of a number of

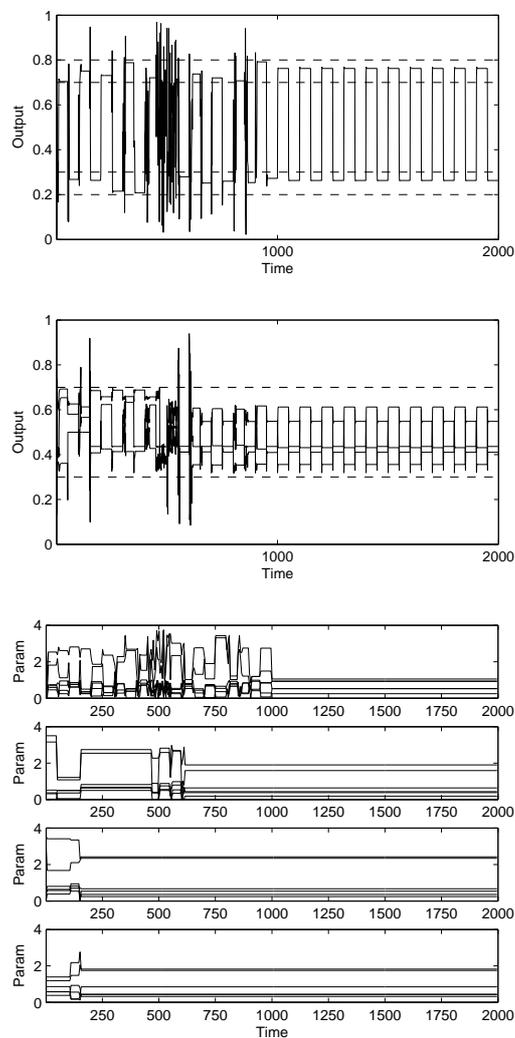


Figure 2.11: 4-unit Homeostat with adaptive transfer function and a designated output unit. When external input to all nodes is low, the output signal should be in the range $[0.2, 0.3]$. When the input signal is high, the output signal should be in the range $[0.7, 0.8]$. External input changes from every 50 timesteps. Eventually a parameter set is found that satisfies the control problem. *Upper*: output signal from designated unit (solid line) and homeostatic target ranges (dashed lines) of $[0.2, 0.3]$ and $[0.7, 0.8]$. *Middle*: output from other units (solid lines) and homeostatic target range (dashed line) of $[0.4, 0.6]$. *Lower*: transfer function (x,y) parameter change.

units arranged in a fully connected network topology. Adaptation occurs by random assignment of connection weights or by changes in transfer function parameters, whenever unit activation goes out of bounds.

The simulated Homeostat was shown to be ultrastable in that a loss of homeostasis in any of its nodes triggered parameter change that in most cases led to a restoration of homeostasis. Stability of the simulated Homeostat required all nodes to maintain homeostasis. The Homeostat was shown to be able to adapt to a variety of perturbations, and was also briefly tested as a control system with a simple input-output mapping task, where it was found able to perform the task when the input signal was experienced directly by the output unit, but unable to perform the task when input was applied elsewhere in the network so that signal transmission was required. This is explained by the homeostatic nature of the units in the Homeostat, which become barriers to signal transmission after adaptation to maintain a steady level of activation. This agrees with Ashby's own observation that homeostatic units are regulators that reduce the effect of disturbance [7].

An interesting observation is that we have observed instances of the simulated Homeostat where homeostasis was inevitable, impossible, or dependent on circumstances. For instance, in Figure 2.6 we can see that one of the nodes can never display homeostasis, since its transfer function always gives output outside the target range and is not alterable by the plastic mechanism. Although an example is not shown, it is easy to think of a case where the transfer function of a node gives output that is always *inside* the target range, i.e., the node is always homeostatic. More common than either of these extremes will be cases where a node can display homeostasis depending on the behaviour of other nodes. This distinction, between a machine that is always homeostatic, a machine that is never homeostatic, and a machine that is sometimes homeostatic, is an important one for homeostatic systems in general.

A related (and better developed) attempt to incorporate homeostatic adaptation and ultrastability into a control system is the body of work presented by Di Paolo [46, 48] in which homeostatic adaptive mechanisms are introduced into controllers for artificial autonomous agents. Some of this work will be covered in the discussion of neural homeostasis presented in Chapter 3, but one aspect of Di Paolo's studies that is of particular relevance here is his exploration of robot controllers that are based on an architecture very similar to the simulated Homeostat used here.

Di Paolo motivates his research by drawing attention to the continuing failure to create artificial agents that can truly be described as intelligent or autonomous. He locates the central problem in the lack of true intentionality in agents created by current methodologies such as evolutionary or behaviour-based robotics, and claims that while these agents may be biologically inspired in some aspects of their structure, they ignore the basic qual-

ities of what it is to be alive. Put simply, because these robots are not alive, they are denied the sense of purpose and the value system that stems from the need to survive. They can have no goals or subjectivity, and any impression of these properties lies in the eye of the observer.

Di Paolo constructs an approach to the design of artificial agents that is grounded in the ideas of Ashby, with a generous nod to the philosophy of Hans Jonas [94] and the concept of autopoiesis [128,129]. Claiming survival as the ‘mother-value’, Di Paolo illustrates the new approach with a robot controller in which sensory input is transduced to motor output via piecewise linear transfer functions similar to those used in the simulated Homeostat described above. Di Paolo’s robotic agent has a battery which is charged by successful performance of phototaxis. The level of charge of the battery forms the essential variable in which homeostasis must be maintained, with plastic change being triggered in the transfer functions of the robot when charge falls below a certain level.

Di Paolo reports successful adaptation of the robot controller, so that phototaxis is reliably performed and battery charge maintained above the threshold level. The robot is also able to adapt to inversion of its sensor array, and even to accommodate periodic inversion and re-inversion. However, Di Paolo notes that there is no guarantee of the time taken to adapt (which can be very long) or that adaptations will be conserved. The simplicity of the task and of the controller also raise questions of scalability. Similar concerns also apply to Di Paolo’s related work on homeostatic adaptation in neural robot controllers (see Chapter 3).

Di Paolo’s work shows that homeostatic adaptation can be successfully applied to controllers for artificial autonomous agents and it is to be hoped that his findings will stimulate further research in this interesting area. There remain significant questions to be answered concerning the ways in which homeostatic adaptation and ultrastability can be implemented in robotic systems, not least in regard to the scalability of such systems. Random parameter search provides powerful validation of the technique, in that it is the worst case scenario for biological adaptation and requires no assumptions concerning mechanism, but it seems likely that more efficient directed search mechanisms will be needed if homeostatic adaptation is to be successfully applied in situations of non-trivial complexity. There is a need for further examples of successful homeostatic adaptation, both to bolster the theory against potential criticisms concerning the simplicity of the tasks currently attempted and to make clear that the successes so far achieved result from the power of a homeostatically adaptive organisation, rather than any unforeseen artefact of implementation.

Ashby’s theory of homeostatic adaptation was developed at a time when the resources needed to push it to its logical conclusions were not available, but in the age of easily

available computational power and increased understanding of biological adaptation, the time may have come for Ashby's ideas to be developed further. Homeostatic adaptation provides an intuitively appealing and widely applicable model for learning, and the idea of ultrastability may be applied in areas other than agent-level adaptation. By exploring the properties of a simulated Homeostat, we hope in this chapter to have contributed towards the foundations of such a study.

Part III

**Homeostatic Plasticity in Neural
Networks**

Chapter 3

Background to Part III

3.1 Overview

Homeostatic mechanisms in the brain have been observed by neuroscientists to act on neuronal and synaptic properties so that neural activity is regulated [37, 167]. Empirical work to further elucidate these mechanisms in biological nervous systems is ongoing. Here we consider the function of homeostasis in artificial neural systems. We look at the effect of homeostatic plastic mechanisms analogous to those observed in biological systems on the dynamics and function of a class of artificial network commonly used in robotics, the continuous-time recurrent neural network (CTRNN) [15].

CTRNNs offer a number of attractive properties but are difficult to parameterise. Currently the best parameterisation methods are based on artificial evolution using genetic algorithms, but this can be problematic and is not yet reliable. Here we examine whether the inclusion of homeostatic plasticity can aid in the consistent evolution of good robot controllers, either by preventing node saturation or by directly improving evolvability in some other way.

Part III of the thesis describes the development and application of homeostatic plastic mechanisms for use in CTRNNs. This chapter gives some background to the research, describing relevant work from the literature and motivating the original research subsequently presented. The remainder of Part III describes experiments intended to explore the effect and utility of homeostatic plastic mechanisms. Chapter 4 looks at how homeostatic plasticity may be incorporated into CTRNNs, before Chapter 5 covers some analysis of

the nature of the homeostatic constraint thus imposed. Chapter 6 describes a set of experiments looking at the generic effects of homeostatic plasticity in large ensembles of randomly generated CTRNNs. Chapter 7 describes some experiments concerning the impact of homeostatic plasticity on evolvability of these networks. Part III concludes with Chapter 8, which offers some discussion of the results achieved and their implications for both robotics and neuroscience.

3.2 Neural and evolutionary robotics

When symbolic artificial intelligence began to founder on the twin rocks of the frame problem [43] and the symbol grounding problem [77, 154], autonomous robotics was forced to look for new approaches. Leading the way were those, led by Rodney Brooks, who thought robotics could best make progress by building real robots that operated in a real environment. In a body of work performed during the late 1980s and 1990s, Brooks presented a framework for ‘new AI’ that emphasised *embodiment* and *embeddness* (see [22, 23] for an overview of this approach). Brooks’ call for robotics to walk before it ran and seek to develop insect-level intelligence led to an increased focus on biologically inspired control. With this began the second coming of neural networks in artificial intelligence.

Neural networks had been around from as early as the 1940s [131] and had been the subject of much research in the 1950s. However, the publication in 1969 of Marvin Minsky’s damning critique [133] condemned them to almost two decades of unfunded obscurity, notwithstanding that the problems Minsky cited were later solved [144, 145]. John Hopfield’s work on memory storage by attractors in recurrent neural networks [86] and Rumelhart & McClelland’s two-volume book on parallel distributed processing [146] heralded a new approach to neural networks. The new style of networks were different to the feed-forward architectures and perceptrons familiar to Minsky. Now neural networks were fully connected and recurrent, allowing cyclical dynamics and state-holding. Later on, more biologically plausible networks with continuous neural activations began to appear.

The new networks were harder to parameterise than feed-forward networks. Normal training methods (such as back-propagation of errors) were unsuitable for the recurrent architecture and new techniques were needed. The increased focus on biologically inspired structures for control led to biologically inspired approaches to parameterisation, i.e., evolution. Genetic algorithms had first appeared in the 1970s [85]. In the 1990s they began to be applied to the optimisation of neural networks for robot control [31, 32, 82] and the field of evolutionary robotics was born.

Evolutionary robotics encompasses a variety of different approaches to robotics that typically use a methodology involving the artificial evolution of neural networks for robot control. Within evolutionary robotics there are many differences in approach, concerning, for example, simulation versus real robots, neural architecture, type of genetic algorithm, etc. Although biological inspiration is widely proclaimed by the community, the degree to which this is actually taken up depends on the particular aims and restrictions of individual projects.

Some researchers want to practice computational neuroethology [13,34], that is, learn about biological nervous systems by looking at the whole-system behaviours of artificial networks. Others simply want to build useful robots, and are not much concerned with biology except as a source of solutions to particular engineering problems. Brooks' watch-words of embodiment and embeddedness are still repeated as mantra by much of the evolutionary robotics community, despite the fact that Brooks' work is neither evolutionary nor based on neural architectures (except in the most general sense of distributed control).

Perhaps the biggest lures of evolutionary robotics are its potential for shedding light on real biological systems and the potential it is thought to offer for automated design. Biological evolution has come up with a rich abundance of examples of robust intelligent control. Perhaps artificial evolution and biological design principles can lead us to robust and intelligent robots.

3.3 Evolving robot controllers

Whatever type of controller is used, the methodology for evolutionary robotics takes a similar form. This involves using a genetic algorithm [68, 85, 134] to optimise the parameters of the controller for an artificial agent so that it performs a particular task. This section will briefly describe the methodology, starting with a general overview of genetic algorithms and then looking at some special considerations when evolving robot controllers.

3.3.1 Genetic algorithms

Genetic algorithms (GAs) are most easily explained by following the steps of a typical example. Pseudocode for a typical GA is given below.

1. Randomly generate an initial population of solutions
2. Encode solutions as a population of genotypes

3. Loop until termination criterion reached:
 - (a) For each genotype in population:
 - i. Instantiate genotype as phenotype solution
 - ii. Test solution on target problem
 - iii. Assign genotype fitness based on phenotype performance
 - (b) Until a new population has been filled:
 - i. Choose two parents using fitness-proportionate selection
 - ii. Combine parent genotypes to form child genotype
 - iii. Mutate child genotype
 - iv. Add child to new population
 - (c) Replace old population with new population.
4. Final solution is best solution from final population

Evolutionary change in GAs occurs by a process of selection between different phenotypes (solutions), coupled with mutation and recombination of genotypes. The form of the candidate solutions depends on the particular problem faced and is not restricted except that solutions must be suitable for encoding as genotypes. The termination criterion for the main loop (Step 3) is often a fixed number of iterations, but may also be the attainment of a fixed level of performance. Step 3a is the fitness testing stage. Step 3b creates a new population of genotypes from the fittest members of the old population. Step 3(b)i is the selection of fit parents from the old population, which are combined using crossover in Step 3(b)ii. Crossover allows recombination of existing solutions, while the mutation of the child genotype in Step 3(b)iii adds new variation to the population. The final output of the GA is the best solution found, i.e., the phenotypic expression of the fittest genotype in the final population.

It should be remembered that the intricacies of GAs are not completely understood and their use is still seen to be something of a 'black art' by practitioners. There is no universal best practice and it is likely that the best type of GA to use is highly contingent on the particular nature of the optimisation problem. Full discussion of all the issues involved with evolutionary computation is beyond the scope of this thesis. However, some of the main stages of the GA are discussed briefly below, while the features of the particular GAs used in this thesis will be given in the text where appropriate.

3.3.1.1 Encoding

Genetic algorithms work by analogy with the genetic code in biological organisms. Candidate solutions to a problem are encoded as an ordered string of symbols (e.g., binary digits) called the genotype. The genotype is a convenient form for conducting the operations of crossover and mutation, which allow new candidate solutions to be created. The new genotypes are then interpreted (via some set of developmental rules) as phenotypes that are candidate solutions to the problem in order to measure their performance.

Genotypes may be encoded as strings of digits that are binary, real-valued or symbolic. Which is used will have an effect on evolutionary dynamics, since it determines the number and effect of possible single-locus mutations. Also, the rules for the genotype-to-phenotype mapping may vary in their complexity and may involve several steps (see for example the grammar re-write rules used by [100]). Genotype-to-phenotype mappings may be one-to-one or many-to-one, since they are generally chosen so that each genotype corresponds to a single phenotype.

3.3.1.2 Crossover

Recombination (crossover) involves the creation of a single child genotype from two parent genotypes during sexual reproduction. Crossover is usually carried out by picking at random a small number of crossover points along the genotype, then switching the parent supplying genetic material at these points. For example, with two parents A and B, single-point crossover in a genotype with 10 loci might be performed by randomly picking locus 7 as the crossover point. The child genotype would then be copied from parent A up until locus 7, at which point the remainder of the genotype would be taken from parent B.

Crossover is not always used. The utility of sexual recombination in GAs is under discussion in the literature and many people use asexual GAs. In an asexual GA child genotypes are created from a single parent and mutation is the only source of genetic variation.

3.3.1.3 Mutation

Mutation involves a random change in value at one or more loci of a genotype. There is a variety of different mutation operators, the most common of which fall into two main categories: point mutation operators and vector mutation operators.

Point mutation can be used with any type of encoding and is carried out by randomly selecting a small number of loci and assigning to them new values (either by adding/subtracting a small amount to/from the existing value or by choosing a new value from a permitted range).

Vector mutation is generally only used with real-valued genotype encodings. It involves adding a vector of values along the length of the complete genotype. Thus it would be too destructive with binary encodings, while addition is generally not defined on arbitrary symbolic encodings. With real-valued encodings vector mutation is typically performed by picking a random point on the M -dimensional hypersphere (where M is the number of genotype loci) to give a unit direction vector, then randomly picking a small vector magnitude from some (typically Gaussian) distribution, before finally adding this vector to the genotype to be mutated.

Point mutation is simple to implement and can be applied with all kinds of genotype encoding, but it mostly creates orthogonal movement through parameter space where only one parameter changes during a mutation event. Vector mutation on the other hand, while less widely applicable, allows movement in any direction through parameter space in a single mutation event by potentially altering all loci at once. The relative merits of each kind of mutation operator have yet to be fully understood.

3.3.1.4 Selection

Many different means of selection have been used in the literature and there is no utility in a review of these here. The reader is referred to [68, 85, 134] for good coverage of a variety of selection operators. Here we will concentrate on the purpose and general form of selection in GAs.

Selection operators use the fitness scores assigned to different phenotype solutions to decide which genotypes should supply genetic material for the next generation, with the idea being that the fittest phenotypes should be most strongly represented. This may be achieved in a variety of ways, but typically some stochastic sampling method is used to favour the selection of fitter parents. One example method is to rank the phenotypes of the previous generation in order of fitness and then use roulette wheel selection. Roulette wheel selection takes the analogy of a roulette wheel in a casino and operates by the random selection of a single slot from all of the possibilities. In a population of size N , roulette wheel selection might give N slots to the fittest genotype, $N-1$ slots to the second fittest, and so on, so that the chance of selection is directly proportional to ranked fitness.

Selection operators are usually associated with generational GAs, i.e., those where the evolutionary search is organised into discrete sequential generations of solutions. The alternative to this is a steady-state GA, which involves a persistent population where individuals are removed or introduced. In this kind of GA selection may be implicit in the ecology of the population, with survival and reproduction taking the place of explicit selection operators choosing parents for the next generation.

3.3.1.5 Fitness function

One important feature of a GA is the fitness function, which determines how credit is assigned to different solutions based on their performance. The fitness of a genotype is a function of some metric(s) measured on the performance of the associated phenotype during a fitness trial. Successful evolution of good solutions relies on a well-designed fitness function. It is not always easy to quantify good performance in terms that can be easily converted into a single fitness score, and often multi-variate fitness scores are used to reward different aspects of performance. Also, scores may be averaged over several trials to reduce noise. However, the end result of any fitness trial is always a score that can be used to compare the worth of different solutions.

The fitness function defines the shape of the *fitness landscape* and hence has a massive impact on the solutions generated by a GA. The fitness landscape is a metaphor that likens the fitness score achieved by a solution based on a particular point in parameter space to the height above sea level in some imaginary mountain range; evolutionary search is then a form of hill-climbing. It can easily be visualised in a two-dimensional parameter space by thinking of fitness as the third dimension.

Two-dimensional parameter spaces are the exception though and most GAs have a large number of parameters. This makes visualisation of the fitness landscape difficult and challenges the utility of the fitness landscape as a good metaphor for evolutionary search. There is some uncertainty over how far it is sensible to regard evolutionary optimisation as hill-climbing, but this is a topic in itself and will not be further explored here; we will tentatively accept some similarity between asexual evolutionary search and hill-climbing on a landscape.

A feature of fitness landscapes that affects the evolutionary search is the ruggedness of the landscape, or its landscape correlation. A highly correlated landscape means that points that are close together in parameter space will have similar fitness. This makes hill-climbing easier, since there is a simple gradient to follow towards well-defined fitness peaks. In a completely non-correlated landscape there is no relation between the distance between two points in parameter space and the similarity of their fitness values. This makes hill-climbing difficult since fitness peaks are more likely to be isolated spikes than smooth gradients, and evolutionary search may become more akin to random sampling than gradient-following.

Another feature of a fitness landscape with relevance to evolutionary search is the amount of neutrality present. Neutrality is the property of adjacent points in the parameter space (in terms of being reachable by a single mutation) having equivalent fitness. Neutrality allows for an amount of genetic drift in an evolutionary search, where new

genotypes may arise without any improvement in fitness. This may be useful for reaching parts of the search space that would not otherwise have been reachable if strict gradient-following was enforced. Neutrality is an open topic in evolutionary theory; see [9] for a good introduction to the topic in artificial evolutionary algorithms.

3.3.2 Using GAs for evolutionary robotics

There are also a number of considerations over GA methodology that are unique to evolutionary robotics. For instance, there is a significant decision to be taken over whether the fitness of a controller is tested by implementing the candidate solution in a real robot or in a simulated robot and environment. Simulation offers a large pay-off in terms of the time taken to perform fitness trials, since working with real robots is time-consuming and fraught with engineering difficulties, but there can be significant problems in transferring evolved solutions from the simulated environment to the real robot. The simplifications necessary in any simulation (such as restricted physics or ‘perfect’ sensors and motors) and the possibility of artefacts in the coded implementation mean that controllers evolved in the simulated environment may not perform well when implemented in the real environment; they have difficulty crossing the ‘reality gap’ [92].

One solution to this problem is to include sufficient noise at critical parts of the simulation to prevent the GA from developing solutions that rely too heavily on any precise feature of the simulation and are thus more likely to be able to cope with changed circumstances in the real world. This approach is known as the ‘minimal simulation’ approach and has been shown to improve the transfer of solutions from simulation to reality [89–91]. Evolutionary robotics also offers some limited scope for evolving the physical structure of agents in addition to, or conjunction with, the evolution of the controller. This is currently easier achieved in simulation than hardware (see [33] for a good introductory study of co-evolving morphology with control in simulation), but various studies suggest that such a process is plausible in hardware [21, 58, 139].

3.4 Continuous-time recurrent neural networks

Continuous-time recurrent neural networks (CTRNNs) [15] are a variety of neural network popular in the evolutionary robotics community for their robustness and general applicability to dynamic control tasks. They offer nonlinear dynamics and have been shown capable of approximating the output of any dynamical system if correctly parameterised [65], meaning that they are suitable for producing the kinds of autonomous and oscillatory dynamics thought to be important for robotic control. They have been used

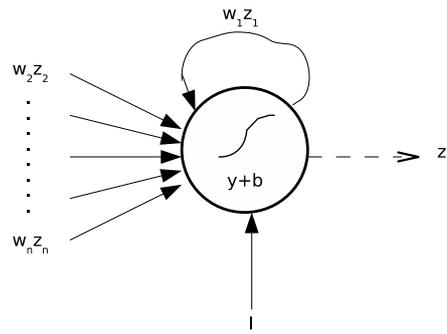


Figure 3.1: Schematic of a single node in a CTRNN network. The node receives input from other nodes ($w_i z_i$), from itself ($w_1 z_1$), and also external input (I). These quantities are summed and contribute to node activation (y), which changes according to Equation 3.1. Node firing rate (z_1) is a sigmoidal function of activation (y) and bias (b), as specified by Equation 3.2.

for robotic tasks including legged locomotion [18, 66], swimming [88] and visual shape discrimination [81], as well as more abstract tasks such as sequence learning [192] and the production of “minimally cognitive behaviour” [10, 155].

This section describes the mathematical formulation of CTRNNs, how they can be evolved using genetic algorithms, and some variants of the standard formulation developed to try and improve performance. The section concludes with discussion of a problematic feature of CTRNNs, that of node saturation, that may be solved by the inclusion of homeostatic plasticity.

3.4.1 Mathematical formulation of CTRNNs

CTRNNs are specified as a set of differential equations that govern how the state of each neuron changes over time and how neuron potential determines firing rate. These are given in Equations 3.1&3.2 below. Figure 3.1 shows a schematic of a single CTRNN node in a network.

$$\tau_y \dot{y} = -y + \sum_{i=1}^N w_i z_i + I \quad (3.1)$$

$$z = \frac{1}{1 + e^{-(y+b)}} \quad (3.2)$$

By analogy with biological neurons, Equations 3.1&3.2 represent the state of a node connected to N nodes including itself, where y represents neuron potential, w_i is the strength of the synapse from the i^{th} afferent neuron, z_i is the firing rate of the i^{th} affer-

ent neuron, I is any external input the neuron receives, and b is the bias term for the neuron. Equation 3.1 defines the rate of change of potential with respect to time (\dot{y}) moderated by a neuron specific time constant (τ_y). Equation 3.2 specifies neuron firing rate as a sigmoid function of neuron potential and bias. Weights can take positive or negative values, representing excitatory and inhibitory synapses. Biases can also be positive or negative, reflecting the neuron's inherent tendency towards quiescence or excitation.

3.4.2 Evolving CTRNNs

Currently the best method for training CTRNNs is artificial evolution using genetic algorithms. The GAs are generally used as previously described; [81] and [10] offer good descriptions of this methodology applied to CTRNNs. Here we will just briefly describe how CTRNNs may be encoded for use with GAs.

Each node in a fully connected CTRNN has $N + 2$ parameters: a decay constant τ_y , a bias term b , and N afferent connection weights w_i from itself and from all other nodes in the network. Thus an N -node CTRNN has $N(N + 2)$ parameters in total. Note that the space of all fully connected N -node architectures contains all N -node networks with lower connectivity, where a connection weight of zero indicates a lack of connection.

Since all the parameters in a CTRNN are real-valued, it is common to use a real-valued encoding. This is the method used for the evolutionary experiments in this thesis. The ranges for connection weights, bias terms and decays are all different, whereas the values used in a genotype usually all come from a single range, for ease of implementation of mutation and crossover. For this reason, CTRNNs are here encoded into a genotype where all values are drawn from the range $[-1.00, 1.00]$, with a developmental stage mapping genotype values linearly to appropriate ranges for their associated phenotypic traits. Further details of the GA and encoding used here will be given in the text as appropriate.

3.4.3 Improving evolvability

Successful evolution of good neural controllers is not necessarily easily achieved and over the years attempts have been made to improve the method in a number of ways, e.g., improvements to the evolutionary algorithm [33, 78, 79], different encoding schemes [64, 72, 100, 132], the addition of Hebbian plasticity [60, 62, 63], the inclusion of mechanisms analogous to the diffusion of gases in biological brains [87], the enforcement of the centre-crossing condition [127], and the use of spiking neurons [50, 59]. These attempts have met with some success, but have not solved the general problem of how to reliably evolve good neurocontrollers. Even where evolvability has been improved, the difficulty

inherent in the post-hoc analysis of such complicated systems means that the reasons for the improvement may not be fully understood.

The type of controller used defines the evolutionary substrate on which the genetic algorithm acts. The different levels of success achieved by different controller types demonstrate that some evolutionary substrates are better than others for certain problems. It is tempting to wonder whether a type of controller exists that is universally better than other neurocontroller variants. However, this kind of thinking is likely to be wrong-headed. An appeal to the principle of ‘no free lunch’ [188, 189] reminds us that there is no universal best optimisation algorithm over all classes of problem. Similarly, there is not necessarily any reason to believe that a universal best neural controller exists for robotics.

However, the lack of a universal best controller does not mean that some types of controller might not be better than other types for a particular class of problems. Also, the range of problems where a particular controller type has an advantage might be broad enough to encompass a large part of the domain of evolutionary robotics. It is this hope that has motivated previous attempts to improve the methodology of evolutionary robotics by developing types of controller that are more evolvable. The mythical substrate of perfectly evolvable, perfectly capable, perfectly robust controllers does not exist, but that does not mean that we cannot try to find controllers that are more evolvable, more capable and more robust than others.

In Section 6.2.6.1 we will argue that any good robot controller will display the basic properties of being able to react to its environment whilst maintaining some internal (possibly oscillatory) dynamics. CTRNNs are capable of providing all these desirable properties; they hold internal state, can propagate a signal from sensors to effectors, and have often been used as central pattern generators (neural circuits that generate autonomous oscillatory dynamics). Their capabilities of universal function approximation and nonlinearity offer further benefits. However, the fact that CTRNNs *can* display these properties does not necessarily mean that they *will*.

3.4.4 The saturation problem

In Chapter 6 we will present experimental evidence to show that randomly parameterised CTRNNs are typically poor at showing the sorts of desirable property we have outlined above. The reason for this is node saturation.

Consider the sigmoidal transfer function that is used to calculate neuron firing rate from neuron potential (given by Equation 3.2). This curve is shown in Figure 3.2, which plots neuron firing rate as a function of potential. Since neuron firing rate is bounded between 0 and 1, in a network of fixed connectivity the effective input received by any

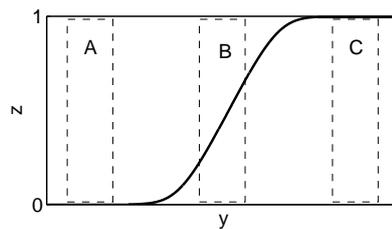


Figure 3.2: Sigmoidal transfer function: firing rate as a function of potential, showing different possible ranges of habitual fluctuation. Ranges A and C give a saturated response where firing rate will not vary significantly in response to a change in input, while range B gives a non-saturated response where firing rate will change significantly.

neuron is also bounded, with the upper and lower limits determined by the number of afferent connections and their weights. This means that the potential of the neuron (being a function of itself and the received input) will fluctuate within a bounded range. The location and size of this range of potential determine what variation is possible in firing rate.

Figure 3.2 shows three examples of different ranges in which potential might vary. Two of these ranges, A and C, give what might be termed a *saturated* response. In these ranges the firing rate of the neuron does not vary despite the variation in potential since the potential is either too high or too low; the neuron is *saturated-off* (range A) or *saturated-on* (range C). The saturated-off response corresponds to quiescence in biological neurons, where the neuron never fires, while the saturated-on response corresponds to biological hyper-excitation, where the neuron always fires at its maximal rate. The third range, range B, gives a non-saturated response. Here the potential varies in the region where the sigmoid curve has the largest gradient, meaning that a change in potential is translated into a significant change in firing rate.

Consider the effect of a change in input on neurons with potential varying in ranges A, B and C. Non-saturated neurons (range B) will change their output when input changes, since the change in potential is converted into a change in firing rate. However, saturated neurons (ranges A and C) will not alter their firing activity, meaning that they display no effective change in state. Saturated neurons thus play no part in network dynamics since they give constant output irrespective of changes in input. They cannot play any part in oscillatory dynamics and act as barriers to the propagation of signals.

Saturation is a continuous, not a discrete, quantity; a neuron may be more or less saturated. It may be thought of in an inverse relation to sensitivity. Fully saturated neurons, where a change in input causes no detectable change in firing, are completely insensitive. Minimally saturated neurons, where the range of potential is centred on the region

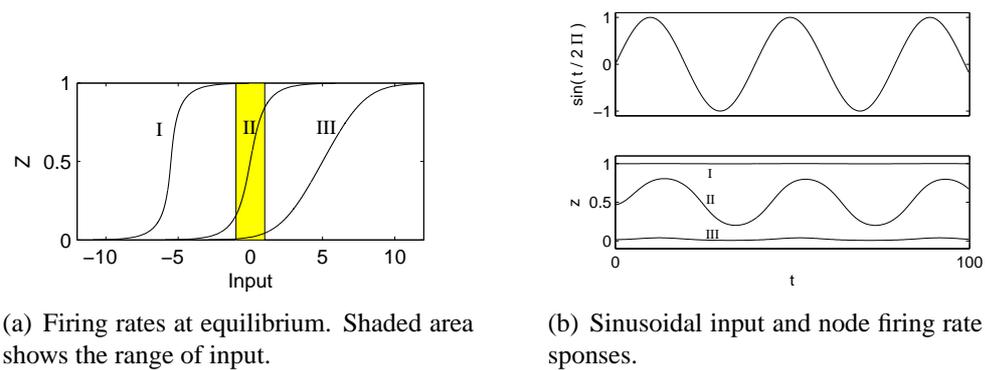


Figure 3.3: Firing rates at equilibrium and firing rate responses to a sinusoidal input signal for three example nodes. Node I has a fully saturated-on response; its firing rate never changes but remains maximal. Node III has a mostly saturated-off response; its firing rate changes slightly. Node II has a strong non-saturated response and its firing rate is changed significantly by the input signal.

of largest gradient of the sigmoid and a change in input causes the largest possible change in firing, are maximally sensitive. Figure 3.3 shows the effect of a sinusoidal input signal on three nodes with varying levels of saturation. The nodes show different levels of response to the same stimulus; one node shows a strong oscillatory dynamic as a result of the external forcing, another a faint oscillation, while another shows no change in activity at all.

Networks will generally be made up of neurons with different levels of saturation, giving an aggregate saturation at the network level. Networks will be more or less saturated/sensitive depending on the saturation levels of their constituent neurons. Networks made up of maximally sensitive neurons are referred to by Mathayomchan & Beer (2002) as centre-crossing networks, where the nullclines of the nodes intersect at their exact centres of symmetry (in other words, where there is an equilibrium point of the system in the absence of input where the firing rate of all nodes is 0.5). Centre-crossing networks were shown by Mathayomchan and Beer to be a fertile substrate for the evolution of rhythmic behaviour, since all nodes played a part in network dynamics. Contrast this with the behaviour of a network made up of fully saturated nodes; such a network would not change its activity in response to stimuli and would not produce interesting behaviours. Depending on the topology and connectivity of the network, networks could be severely affected by saturation of just a few nodes at key locations. Centre-crossing networks may be thought of as *poised* to behave, while saturated networks may be thought of as *inert*. We can expect signals and waves of activation to propagate well in poised networks, but not at all in networks that are inert.

3.5 Homeostatic plasticity in biological nervous systems

It is useful to consider biological neural networks and to observe that saturation effects (i.e., hyper-excitation and quiescence) are probably not a common problem. One postulated reason for this is the existence of homeostatic plastic mechanisms that serve to regulate neural activity [37, 167]. While the precise feature of neural activity that is regulated is not known (it may be mean firing rate, mean calcium concentration or some other feature) it is clear that neural activity tends towards a constant level in the long term. It is also clear that there are a variety of mechanisms by which this homeostasis is accomplished, amongst which are a number of mechanisms affecting the strength of synaptic connections [1, 25, 37, 167] and several mechanisms affecting the intrinsic excitability of individual neurons [2, 45, 167, 194].

Biological neuronal networks are highly complex systems. As well as the electrical firing of neurons and release of neurotransmitters, the structure of the network and internal properties of neurons all change continually. Experience-dependent plasticity and developmental change cause alterations to the network, yet in the face of this continual perturbation the network somehow maintains a large amount of stability. Homeostatic plasticity acts to promote stability in neuron firing rates. It is thought that homeostatic plasticity may counteract the positive feedback effects associated with Hebbian learning [171] and that it may also serve to maintain stability in neuronal networks during development [172].

This section will give a brief description of some of the forms of homeostatic plasticity observed by neuroscientists. The literature on this subject is very large and a complete review is beyond the scope of this thesis. The interested reader is referred to [37, 167, 171, 172], which together give a reasonable overview of current knowledge of homeostatic plasticity in biological nervous systems. In brief, neurons can use their own activity as a feedback signal that allows them to maintain long-term stability in firing rate by homeostatic mechanisms. There is a wide variety of mechanisms, of which only a few will be described here. For our purposes we are mainly interested in the mechanisms that are most suitable for use with artificial neural networks, which include those affecting the intrinsic excitability of neurons and those regulating neuronal activity by multiplicative scaling of afferent synaptic connections. We will describe these below, together with a brief description of a few other kinds of homeostatic plastic mechanism which are included to indicate the variety of mechanisms so far observed.

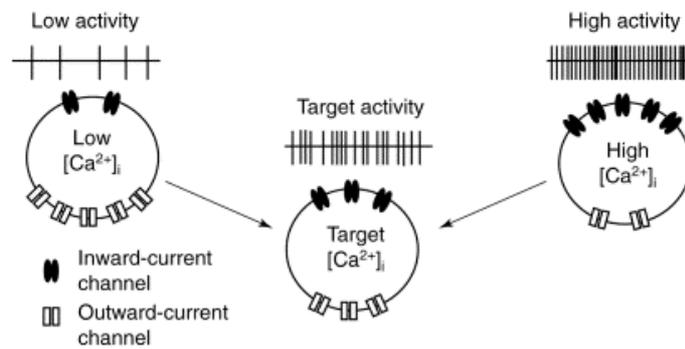


Figure 3.4: When firing rate is low, intracellular Ca^{2+} concentration falls, triggering a compensatory increase in the inward current and decrease in outward current. This raises the excitability of the neuron, increasing firing rate. The opposite occurs when firing is too high, increased Ca^{2+} concentration leads to a reduction in inward current and an increase in outward current. Excitability is then lowered, reducing the firing rate. In this way, intracellular Ca^{2+} concentration is homeostatically regulated around a target level, with an associated regulation of firing rate. Diagram adapted from [167] (p.222).

3.5.1 Plasticity of intrinsic excitability of neurons

The levels of synaptic input to a neuron may vary dramatically during learning and development, yet neurons somehow remain sensitive to input, suggesting that their output response is regulated dependent on their recent history of activity. The integrative properties of a postsynaptic neuron are determined by the mixture and distribution of the voltage-dependent sodium, calcium and potassium conductances that it exhibits. Thus if activity can alter these conductances, it can alter the response of the neuron to input, that is, it will alter the intrinsic excitability of the neuron [171]. Such an effect has been demonstrated in invertebrate neurons, where ongoing patterned activity was shown to regulate conductances [69, 70, 168, 170]. A similar effect has been observed in cultured neocortical pyramidal neurons, where activity blockade lowers the threshold for spike generation, raising the firing frequency for any given level of input stimulation [45].

It is modifications to the balance of inward and outward ion channels that performs the regulation of firing activity. Regulation is triggered by some signal that is well-correlated with firing activity, such as intracellular Ca^{2+} [167]. If the neuronal firing rate is low, the intracellular Ca^{2+} concentration falls, triggering an alteration in conductances that raises firing rate (and with it Ca^{2+} levels). If the firing rate is too high, intracellular Ca^{2+} concentration rises and conductances change to decrease activity and lower Ca^{2+} levels. This is shown in Figure 3.5.1.

Homeostatic regulation of conductances is an example of how the intrinsic excitability of neurons may be altered to maintain long-term stability in firing rate, where the target

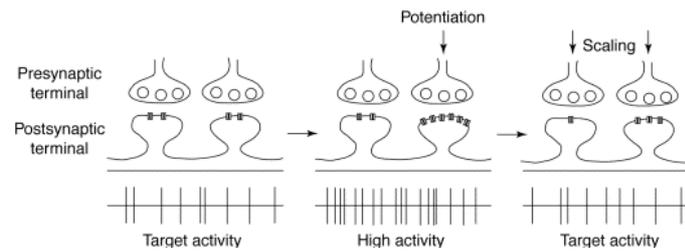


Figure 3.5: Synaptic scaling preserves relative differences between synaptic strengths. Initially, the two synapses have equal strength. When one synapse is strengthened by long-term potentiation, postsynaptic firing rate increases, causing a scaling down of the strengths of both the synapses. Thus firing rate is regulated and relative strengths of synapses are preserved. Diagram adapted from [167] (p.224).

level of the Ca^{2+} signal reflects the set-point around which firing rate is regulated. Such regulation may be able to fine-tune the output properties of a neuron to match its input, so that the information encoded by the firing rate output is maximised [158].

3.5.2 Synaptic scaling in the central nervous system

Regulation of neuronal activity in the central nervous system (CNS) is difficult. Each neuron may have thousands of afferent synaptic connections, which may be inhibitory, excitatory or modulatory. The number and strength of these connections also changes during the lifetime of the neuron. This input complexity makes the maintenance of stable activity by a neuron unlikely, but it now seems likely that cortical and hippocampal neurons regulate their own firing by scaling their synaptic inputs up or down as a function of activity [167]. This scaling has been demonstrated in cultured cortical neurons, where excitatory connections between pyramidal neurons were globally scaled in relation to firing rate; when firing rate was high, excitatory connections were scaled down, when firing rate was low excitatory connections were scaled up [169]. The scaling is achieved by changes in the quantal amplitude of the AMPA receptor-mediated component of excitatory neurotransmission, which is responsible for much of the excitatory transmission in the CNS.

Synaptic scaling in the CNS is a slow process, taking hours or days of altered activity to modify synaptic strengths [167]. This allows short-term fluctuations (necessary for encoding information in normal operation) while ensuring that long-term stability is maintained. Also, synaptic scaling is a multiplicative process that acts equally on all excitatory synapses to a neuron. This preserves relative differences in strength between synapses, such as those caused by synapse-specific plasticity like long-term potentiation or long-

term depression (see Figure 3.5.2). Synaptic scaling can also lead to competition between synapses if (as is believed) there is some regulation of total synaptic strength [167].

3.5.3 Synaptic homeostasis at the neuromuscular junction

One of the earliest demonstrated examples of homeostatic plasticity is the response to loss of innervation to skeletal muscles. When synaptic drive is lost, muscles become more excitable and likely to contract spontaneously. Genetic manipulation of synaptic properties in *Drosophila* has highlighted the existence of a number of compensatory mechanisms that act to keep neuromuscular transmission relatively constant [37]. At the *Drosophila* neuromuscular junction there is a very narrow range of innervation where the muscle will respond appropriately, outside of which the muscle either becomes hyper-innervated (causing tetanus) or hypo-innervated (causing failure to contract). However, when one synaptic property is altered by genetic manipulation, other aspects of synaptic transmission change to homeostatically regulate the level of innervation [37, 38].

3.5.4 Neurotrophins and regulation of activity in cortical networks

Brain-derived neurotrophic factor (BDNF) has been implicated in various mechanisms of activity-dependent plasticity, including the modulation of synaptic transmission, long-term potentiation, postsynaptic depolarization, dendritic outgrowth, synaptic scaling, and plasticity of intrinsic neuronal excitability [171]. Some studies show that short-term exposure may increase excitatory synaptic transmission, while others show that longer-term exposure may act to stabilise network activity.

Long-term, low concentration exposure to BDNF appears to stabilise the activity of cortical networks, not only by regulating intrinsic excitability [44], but also by balancing the strength of inhibitory and excitatory inputs [147]. In the cortex, pyramidal neurons have an excitatory effect on other pyramidal neurons, while interneurons have an inhibitory effect. BDNF acts to scale down excitatory inputs to pyramidal neurons and scale up those to interneurons, raising the level of inhibition in the network. Because the level of BDNF is positively correlated with activity, this mechanism acts as a negative feedback on activity that maintains overall levels of firing in the network around some set point.

3.5.5 Activity-dependent regulation of synapse number

Perhaps the most controversial form of homeostatic plasticity in the brain, with evidence in the literature both for and against it, concerns the idea that regulation may be achieved

through changes in the number of synapses and not just by changes in their strength.

Some studies have suggested that numbers of AMPA and NMDA receptors can be modulated as a result of prolonged changes in activity, though whether they are regulated independently or in parallel is unknown [171].

Homeostatic plasticity may play a role in synaptogenesis, since the success or failure of the possible synapses during dendritic/axonal outgrowth and retraction is determined by the strengthening of the synapse [167]. Since homeostatic plasticity affects synaptic strength it must implicitly affect the success of new synapses. It is also possible that there is some global regulation of the number of synapses, since some similar process is needed to counteract Hebbian mechanisms and explain retraction of synapses during development [171]. This is related to the competition between synapses that is fostered by synaptic scaling; synapses weakened by scaling down in response to strengthening of other synapses may reach a point where they are eliminated [167].

3.6 Homeostatic plasticity in artificial nervous systems

There are few previously published studies of homeostatic plasticity in artificial neural networks for robot control. The most significant papers are those by Di Paolo [46,48] on homeostatic adaptation (followed up by Balaam [8]), and another paper by Di Paolo on the use of homeostatic oscillators to give functional robustness [47]. Another related article is [84], which also attempts to apply homeostatic plasticity to artificial neural networks. Also related is the concept of centre-crossing networks, presented in [127]. These papers are briefly reviewed below.

3.6.1 Di Paolo (2000) ‘Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions’

Di Paolo (2000) [46] presented an investigation into Ashby’s ideas of ultrastability and homeostatic adaptation [7] applied to evolutionary robotics. Di Paolo was interested in the link between internal stability and adaptive behaviour, and whether a causal relationship could be established between them in an artificial agent; if this could be done then perhaps it would give similar robustness and adaptability to that seen in animals.

Di Paolo tested his ideas in a simple evolutionary robotics scenario. He evolved CTRNN controllers for a simulated robot to perform phototaxis in a simple environment. Fitness was awarded for performance of phototaxis on a series of light sources, and also for maintenance of internal stability. Internal stability related to a novel feature of the CTRNNs that Di Paolo used, which was the inclusion of a form of neural homeostasis.

Di Paolo combined the notion of node-level homeostasis in firing rate with Hebbian learning mechanisms acting on synapses. Each synapse had a genetically specified Hebbian learning rule associated with it, which was applied when post-synaptic firing was too high or too low. This selective plasticity was implemented through the use of a plastic facilitation function. Plastic facilitation was a piecewise linear function of neural activation that returned a signed value when activation was outside a target homeostatic range.

Controllers were evolved with fitness awarded for phototactic behaviour and for keeping node activation inside the homeostatic range. Evolved controllers were then tested for long-term stability by running them for several hundred light source presentations, to ensure that there were no destabilising slow dynamics occurring. The stable controllers were then tested for adaptation to a variety of sensorimotor disruptions, with results presented for adaptation to inversion of the visual field.

The simulated robot had two light sensors well separated on a circular body (i.e., the same layout as the robot shown later in Figure 6.1). Inversion of the visual field in this scenario corresponds to switching the input to the left and right sensors. This switch was performed and then robot behaviour was observed to see if adaptation (recovery of phototactic behaviour) occurred. Di Paolo's hypothesis was that phototactic behaviour and internal stability (node homeostasis) somehow become linked during evolution so that they require each other. When a perturbation occurs that disrupts the phototactic behaviour, the resulting loss of internal stability will drive plastic change until a new stable phototactic attractor is discovered. This should occur since internal stability can only be recovered when phototaxis is performed; the controller is ultrastable, similarly to the Homeostat [7].

Results showed that adaptation to inversion occurred in around 50% of stable controllers. The time taken to re-adapt was linearly correlated with the time of inversion, i.e., it took longer to adapt when the sensors were swapped later in the agent lifetime. Similar adaptation was reported to other perturbations, such as asymmetric alterations to the gain of sensor and motor neurons.

The work presented by Di Paolo in [46] is interesting and is (to the best of the author's knowledge) the first reported attempt to incorporate any form of neural homeostasis into an artificial neural network for robot control. However, there are a few areas of uncertainty over the results. Firstly, the simplicity of the task makes it hard to be sure that genuine re-adaptation has occurred. Phototaxis can be easily performed by a light-sensitive robot using only one sensor and moving in a cycloidal motion towards the source. While the robot motion reported by Di Paolo prior to inversion is a smooth two-sensor solution, the motion by re-adapted agents after the inversion is reported as this kind of cycloidal trajectory. The existence of these single-sensor solutions raises questions over the suitability of

the task for testing adaptation to sensor inversion.

Also, it is not clear how successful the evolved controllers are at maintaining internal homeostasis. The nature of the Hebbian rules used means that positive feedback and runaway weight change can occur, and in these cases non-homeostatic saturated attractors will be easily reached, where weights are pushed to the extremes of their permitted range by the plasticity. This kind of feedback effect has been observed by the current author in unreported investigations of Hebbian learning in CTRNNs.

A final concern that is less to do with Di Paolo's work and more related to our purpose in this thesis, is that the homeostatic plastic mechanisms used by Di Paolo are not reflective of the view of homeostatic plasticity from neuroscience. The idea of plasticity triggered when firing goes outside a target range is valid, but the use of Hebbian mechanisms is not. Homeostatic plasticity in biological brains is thought to be a counterpart to Hebbian learning, using node-level scaling and intrinsic plasticity to give negative feedback on activity. Hebbian mechanisms, however, are synapse-specific and tend to give positive feedback; it is precisely these effects that homeostatic plasticity is thought to prevent [171].

Di Paolo's study was subsequently re-implemented by Balaam [8], who performed further experiments and minor modifications to the original scheme. Balaam found that successful re-adaptation occurred less frequently than the 50% rate found by Di Paolo. Balaam speculated that phototaxis might be a behavioural attractor with a large basin of attraction, and measured the likelihood of phototaxis occurring with randomly generated control networks. He found that almost 10% of random networks of the same size as Di Paolo's networks performed phototaxis, and raised concerns that adaptation in this scenario might not be as difficult as previously assumed.

Despite these concerns, Di Paolo's work remains an interesting and original piece of research. Its aim was to study homeostatic adaptation in artificial agents, rather than homeostatic plasticity in neural networks, and it has provided a valuable first step in this line of questioning.

3.6.2 Di Paolo (2003) 'Organismically-inspired robotics: homeostatic adaptation and teleology beyond the closed sensorimotor loop'

In [48], Di Paolo further develops his theoretical model of biological homeostatic adaptation and argues that robotics should move beyond just taking biological inspiration for the solution of engineering problems, and seek to incorporate the fundamental properties of biological organisation into artificial agents in order to give them true autonomy and self-generated purpose. As such, the paper is largely a philosophical treatment of the na-

ture of life and cognition, where homeostatic adaptation is featured in its Ashbyan sense of maintaining organisation and protecting essential variables. Homeostatic plasticity in neural systems is included in the form of a review of Di Paolo (2000), where he highlights similar problems to those mentioned above. He presents a simple model that addresses some of these problems, in the form of a robot controller that performs a piecewise linear input-output mapping and behaves as Ashby's Homeostat [7]. The lack of a neural model in this work means that it does not offer any useful advance for our current purpose; discussion of this work is included here for completeness and to demonstrate Di Paolo's own awareness of the problems with his earlier work that we described above.

3.6.3 Di Paolo (2002) 'Evolving robust robots using fast homeostatic oscillators'

In [47], Di Paolo presents an exploration of the how robustness of evolved neurocontrollers can be improved using neurons that are constrained to oscillate at a timescale faster than that of behavioural dynamics. He showed that robustness to various sensorimotor disruptions was increased because the faster timescale forced behavioural function to be distributed across the whole control system, meaning that each component played a smaller role and hence had a smaller impact on performance if perturbed.

The networks Di Paolo used were variants of CTRNNs where the bias term becomes a variable.¹ The node state and transfer functions are the same as those given by Equations 3.1 and 3.2 respectively, while the bias was now governed by Equation 3.3 below.

$$\tau_b \dot{b} = -(b + y) \quad (3.3)$$

where b and y are the bias term and activation of a neuron as before and τ_b is a genetically set neuron-specific time constant constrained to be a multiple of the decay constant for that neuron. The bias thus acts on a slower timescale than the activation, and drives firing rate to an average of 0.5 in the long term.

Fully connected networks with 4, 6, 8, 10 and 20 nodes were instantiated as controllers for a photosensitive robot (again similar to that shown in Figure 6.1). Controllers were evolved to perform phototaxis in a simple environment containing a light source, then the evolved controllers were tested for their robustness to various sensorimotor disruptions.

Results showed that the controllers with the adapting bias were more robust, retaining a significantly greater proportion of the unperturbed performance when sensors were swapped, lesioned or suffered angular displacement. Analysis showed that the nodes in

¹Note that bias adaptation is sometimes referred to as threshold adaptation in the neuroscience literature.

the evolved controllers all performed as oscillators. Di Paolo’s conclusion was that this prevented individual nodes from taking on functionally specific roles in the performance of a behaviour, distributing behavioural control across the whole network so that failure of individual elements had a lesser effect. He ended the paper with a hypothesis, “In complex multi-component systems, robustness will be likely to be obtained if functional specification at the level of individual components is minimised” (p.8, [47]).

While Di Paolo’s work with homeostatic oscillators makes use of a homeostatic plastic mechanism that is more directly related to one of the mechanisms observed by neuroscientists (plasticity of intrinsic excitability), his main emphasis was on improving robustness of evolved controllers by enforcing multiple timescales. He did not address the questions of the general effects of homeostatic plasticity on neural network controllers, did not look at how it affected evolvability, and only considered one type of mechanism. The work is valid and relevant to this thesis, but leaves substantial gaps in understanding that we hope will be addressed by the original work presented here in later chapters.

3.6.4 Mathayomchan & Beer (2002) ‘Centre-crossing recurrent neural networks for the evolution of rhythmic behaviour’

In an earlier paper [15], Randall Beer had speculated that centre-crossing networks might be more evolvable. Mathayomchan and Beer (2002) [127] tested this hypothesis and found that seeding evolutionary searches with centre-crossing networks led to quicker evolution and better solutions.

To quote Mathayomchan and Beer directly, “A center-crossing CTRNN is one in which the null-(hyper)surfaces of individual neurons cross one another at their exact centers, ensuring that the range of inputs that each neuron receives is centered over the most sensitive part of its activation function” (p.2044, [127]). This means that there is a stable fixed-point equilibrium state of the network where the firing rate of each neuron is 0.5 exactly. Substituting for this value into Equation 3.1 and solving gives a relation for the bias and weights of a neuron at the centre-crossing condition:

$$b^* = \frac{-\sum_{i=1}^N w_i}{2} \quad (3.4)$$

where b^* is the bias value that will give the centre-crossing condition for the set of afferent weights w_i .

The argument that centre-crossing networks will be more evolvable is based around the observation that, “Due to the form of [Equation 3.2], unless a neuron’s bias is properly tuned to the range of inputs it receives, that neuron will simply saturate on or off and

drop out of the dynamics. Thus the richest dynamics should be found in the neighbourhood of the center-crossing networks in parameter space, and one would expect that an evolutionary algorithm would benefit from focusing its search there.” (p.2046, [127]).

Mathayomchan and Beer tested this hypothesis by seeding an evolutionary search with centre-crossing networks. The task was to evolve CTRNN central pattern generators (CPGs) for walking in a legged robot, a task which had already been well-studied by Beer in previous work [12]. Mathayomchan and Beer compared evolutionary searches seeded with randomly generated CTRNNs and with CTRNNs where the weights were randomly generated but where the bias terms were calculated to give centre-crossing networks using Equation 3.4.

Results showed that the seeded searches consistently evolved higher fitness solutions and that higher fitness solutions evolved faster in seeded searches. Mathayomchan and Beer found that in general, centre-crossing networks were much more likely to display oscillations than ordinary networks, so that the initial population was likely to contain more fit solutions when the centre-crossing condition was enforced. This gave more potential oscillatory circuits to begin with, which could then be refined by the evolutionary algorithm to fit the control task.

While the results given by Mathayomchan and Beer are based on a walking task, they argued that because centre-crossing networks are more suited to producing CPGs, they will be more evolvable for any oscillatory task. They also suggest that centre-crossing networks may be more evolvable in general, since they allow easier access to a wider range of dynamics.

3.6.5 Hoinville & Henaff (2004) ‘Evolving plastic neural controllers stabilized by homeostatic mechanisms for adaptation to a perturbation’

One paper that explicitly tried to incorporate a more biologically plausible model of homeostatic plasticity into neural networks for robot control was that of Hoinville and Henaff [84]. Hoinville and Henaff were interested in increasing the robustness of evolved neurocontrollers to external perturbation. They presented a model that included both Hebbian and homeostatic plasticity, and tested their networks on a simulated single-legged robot. While the intentions of the work presented by Hoinville and Henaff were good, there are a number of problems with their work that mean it may not have achieved its goals.

Hoinville and Henaff used a variant of the standard CTRNN and added Hebbian plasticity rules that they based on the adaptive synapses model presented in [61]. In the adap-

tive synapses model, every synapse has an associated learning rule that changes synapse strength according to some function of the correlated firing of pre-synaptic and post-synaptic neurons. This rule is genetically selected from four possibilities (plain Hebb, pre-synaptic, post-synaptic, covariance), adding a further N^2 symbolic loci to the genetic encoding. The use of this kind of plasticity has been explored in (e.g.) [46, 60, 62–64], which demonstrate some interesting results without a full analytical understanding of how such plasticity affects the network.

To this already complicated network architecture, Hoinville and Henaff then add two forms of homeostatic plasticity. One is regulation of intrinsic neuronal excitability using a rule based on the idea of centre-crossing networks discussed in [127]. The other is normalisation of synaptic weights.

Hoinville and Henaff mistakenly claim that a centre-crossing network can be made by switching from the asymmetrical logistic transfer function (Equation 3.2) to a symmetric function such as the hyperbolic tangent. This claim does not hold.

Normalisation of weights is performed by dividing the magnitude of every afferent synaptic weight to a neuron by the magnitude of the resultant of the vector of afferent weight magnitudes. This is performed after every change to the weights caused by the Hebbian plasticity rules and keeps the size of the resultant equal to 1. This stops runaway weight change and prevents grossly excitatory or inhibitory weightings from occurring. However, it ignores the sign of the weights, meaning that it does not allow for the case where large excitatory connections are balanced by large inhibitory connections.

Despite the problems with each of these rules individually, there is a more significant problem that affects both of the homeostatic rules used by Hoinville and Henaff. This is that neither of the mechanisms that are supposed to ensure homeostasis of firing rate actually take into account the current level of firing. The linear mapping of afferent firing rates from the range $[0, 1]$ to $[-1, 1]$ does nothing to change the current firing rate of a neuron and is based on an incorrect understanding of the centre-crossing condition. The normalisation of weights will regulate the size of the inputs to a neuron to some extent but there is no feedback between the firing of a neuron and the weight normalisation except indirectly through the Hebbian mechanisms.

Hoinville and Henaff showed that controllers using various combinations of the above plastic mechanisms could be evolved to give robust maintenance of a target velocity in a single-leg robot system. They also showed that when perturbations such as changes to friction coefficients or to the target speed were introduced to the evolutionary process the genetic algorithm could still find good solutions. This result is valid and adds another data point to the set of successful evolved neurocontrollers. However, the bigger claim of Hoinville and Henaff, that their ‘homeostatic’ controllers are more robust to external per-

turbations than more standard controllers, is unfounded. Setting aside the problems with the implementation of their homeostatic controller, the fact that the evolved controllers are only shown to be robust to the perturbations which were present in their fitness trials and for which they were evolved means that any claim of general robustness must be watered down to a more modest claim of meeting the fitness criteria.

While the aims of the paper presented by Hoinville and Henaff are good, and the attempted inclusion of different sorts of homeostatic mechanism is to be commended, it is the opinion of this author that there are too many flaws, with both the design of the controller and with the experimental set-up, to be able to learn much from this work. The way in which homeostatic plasticity has been included into the control network should be reconsidered, the evolved controllers should be tested against perturbations that were not present during evolution, and the post-hoc analysis should be much more rigorous if this work is to be useful.

We should no longer be surprised that a controller with so many variables and undoubtedly complex network dynamics can be optimised by a genetic algorithm to give performance of a simple task. Instead we should be asking why and how it does so, and there is simply too much complexity in the networks presented by Hoinville and Henaff to be able to isolate cause and effect. Rather than adding every plastic mechanism we can think of, all at once, we should be seeking to describe the properties of each mechanism and its effect on network dynamics and evolvability, by a methodical series of experiments and analysis.

3.7 Can simple simulations inform neuroscience?

In this chapter so far we have seen a clear flow of ideas from neuroscience to robotics, but it is worthwhile to consider the possibility that there can also be a flow of information in the other direction. Although the neural models used in robotics are hugely simplified and abstracted compared to real brains, there are a number of areas in which it may be possible for neural robotics to inform neural science.

Computational modelling is an accepted methodology in neuroscience, where detailed models of different brain structures are simulated and compared to real-world observation. However, computational neuroscience has a tendency either towards detailed models that are computationally expensive (and may in any case miss out crucial features of the real system), or towards simple models that are tractable but bear no relation to the real system. In attempting to capture the complexity of real neural systems, computational constraints mean that there is always a trade-off between the level of detail in a model and the scale of the system that can be modelled. It is not yet possible to model whole brains at a high

level of neuronal detail, and will not be possible for some time.

CTRNNs are highly abstracted and are probably not very good models of real neuronal networks. However, their abstraction means that relatively large networks can be modelled, and their suitability for use as robot controllers opens up the interesting possibility of modelling at the level of an agent interacting with an environment. This is something that has never before been possible. The scale and complexity of biological neuronal networks, and the computational expense of accurate models, have meant that experiments on behaving neural systems have not been possible until the advent of neural robotics. We can't observe undisturbed biological brains behaving in their natural state, neither can we selectively alter different parts of working brains and measure the effect on animal behaviour. But we can do exactly that in artificial neural systems, where all parameters and variables are accessible, and where experiments are repeatable and free of ethical concerns.

Neural systems used to control agents that act in some environment close the sensorimotor loop; the agent is involved in a continuous reciprocal interaction with the world. As such, sensory information and patterns of neural activity have meaning that is grounded in behaviour. Rather than arbitrarily assigning inputs and meanings, we are now able to observe the actions of an agent in its environment and interpret neural dynamics in terms of objectively measurable behaviours. This grounding avoids the problems inherent in subjective speculations over the possible function of different types of neural mechanism for higher-level functions; we do not need to guess at how neural activity might affect agent-level behaviour because we can directly measure it. The symbol-grounding problem of symbolic AI (and non-embedded connectionist models) is avoided.

The utility of this whole-system analysis for learning about biological brains rests on how well the artificial networks we use share the properties of the biological ones we want to learn about. It is readily apparent that CTRNNs are not accurate models of real neuronal networks; they are highly abstracted and 'neurons' in CTRNNs are probably better interpreted as representing groups of biological neurons. However, it is also apparent that both CTRNNs and biological neuronal networks fall into some higher category of 'parallel distributed processors', or more ambitiously, the general class of 'neural systems'. As such, it may be that learning about one can help us to learn about the other. Real neuronal networks are made up of interconnected nodes displaying a sigmoidal firing response to stimulation, just like CTRNNs. Maybe a good understanding of the dynamics of CTRNNs and behaviours of agents controlled by them will help us to a better functional understanding of animal brains and animal behaviour.

A good example of this kind of work is Beer's analysis of the neural controller of an agent that had been evolved to perform a categorical perception task [11]. In this work,

Beer applied dynamical systems analysis to the coupled network/agent/environment system to understand the way in which a decision was reached. In doing so he highlighted a number of interesting issues. For instance, the ‘decision’ made by the agent over which category an object belonged to appeared to be extended both temporally during an active scanning behaviour and spatially across the whole system. Describing the decision-making process in terms of dynamics and attractors raises interesting questions for cognitive scientists and provides a concrete example for those who take a dynamical systems approach to cognition [14, 16, 164, 173]. Do biological brains follow a similar decision-making process? Might other cognitive tasks have a similar explanation? The answers to these questions are far from clear, but Beer’s approach is surely valid.

In justifying his work, Beer makes reference to the metaphor of the ‘frictionless brain’ [11]. Just as Newtonian mechanics was developed and made useful by ignoring complicating real-world factors such as friction and air resistance and instead working with idealised representations of the world, it may be that by abstracting the fundamental principles of neural systems away from the messy complexity of biological brains we may arrive at a better functional understanding of brain-like systems.

Such a methodology is to some extent followed here. We don’t pretend that the models of homeostatic plasticity presented below are accurate models of how it occurs in biology, but we do maintain that the mechanisms used share some of the fundamental properties that their biological counterparts must have. So long as we remember that we are working with simple abstractions it may be that we, like children, can learn about the real world by playing with toy models.

3.8 Aims, objectives, questions

The over-arching aim of evolutionary robotics is to build good robots, and this must be remembered as the primary goal here. However, to do so is not a straightforward matter of iteration and improvement. There is also room for new ideas and exploratory study. We have seen that CTRNNs and their variants are capable of producing good robot controllers, but also that they do not easily or reliably do so. We have identified a possibly problematic feature of CTRNNs that might make them hard to evolve, i.e., that of node saturation. We have also observed that node saturation is not a problem in biological nervous systems, and that this is likely to be because of the presence of homeostatic plastic mechanisms that act on system parameters to keep activity in bounds.

We feel that these observations justify a study of the utility of homeostatic plasticity in artificial neural networks. Specifically, we want to find answers to the following questions:

1. Can HP be implemented in CTRNNs?
2. Does HP prevent node saturation?
3. What effect does HP have on node/network/agent behaviour?
4. Does HP make CTRNNs more evolvable?

These questions will be described and justified more fully, and hopefully answered, in the remaining chapters of Part III. While the emphasis of Part III is on the engineering aspects of these questions, we will always try to discuss the relevance of the work presented for neuroscience.

Chapter 4

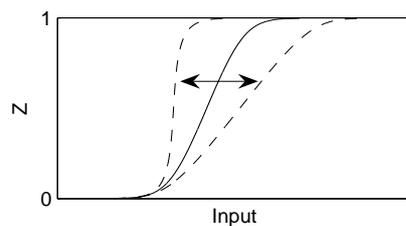
Can homeostatic plasticity be implemented in CTRNNs?

4.1 Overview

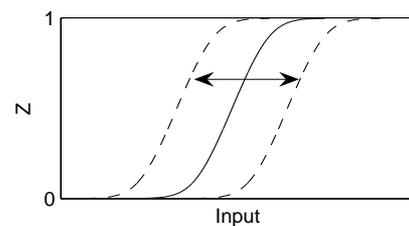
Before any study of the effects and utility of homeostatic plasticity in continuous-time recurrent neural networks (CTRNNs) can be undertaken, we first of all need to determine how it should be implemented. We need to consider which of the various kinds of homeostatic mechanism are suitable for incorporation into CTRNNs and then develop a formulation for how this should be done.

4.2 What types of mechanism should be implemented?

As described in Chapter 3, there are many different sorts of homeostatic plastic mechanism in the brain, and therefore there are a wide variety of candidate mechanisms that we might adapt for use in CTRNNs. With enough effort, almost any of the mechanisms might be used, but given the simplified nature of CTRNNs compared to real brains it seems sensible to work at a similar level of abstraction when adding plasticity. CTRNNs are highly abstracted, highly idealised simulacra of small sections of biological neuronal matter that ignore the full complexity of neural dynamics in favour of a conceptually clear framework for study. The homeostatic plastic mechanisms developed here should follow the same rationale, for reasons of analytical and computational tractability, and also because over-



(a) Synaptic scaling applied to the self-connection affects the gain of a neuron, changing the slope of its input-output function.



(b) Internal plasticity can affect the intrinsic excitability of a neuron, translating its input-output function. Synaptic scaling of input weights has a similar effect.

Figure 4.1: Schematic of the effects of different kinds of homeostatic neural plasticity. All axes are linearly scaled.

complicating the plasticity we add to a CTRNN will move the resulting networks too far from the standard CTRNN for reasonable comparisons to be made.

CTRNNs (as used in the robotics community) typically have a fixed architecture that does not change during the functional lifetime of an agent, i.e., no new synapses are added or removed. Adding homeostatic mechanisms involving the creation or deletion of connections would therefore involve considerable effort; first of all a reasonable model of how neural architecture might change *without* homeostatic constraints would be needed, before a model of how this might occur *with* homeostatic constraints could be developed. For this reason we will not develop mechanisms that change network topology.

Broadly speaking, the remaining homeostatic mechanisms that are well-documented in the literature can be grouped into those that act on synaptic efficacies and those that affect the intrinsic excitability of neurons, and these will be the mechanisms we study in this thesis. Rather than copy individual mechanisms in detail, we will appeal to simplicity and develop mechanisms that hopefully capture the qualitative functional performance of the myriad individual synaptic and internal mechanisms. That is, we will develop two simple homeostatic plastic mechanisms that perform synaptic scaling and that alter intrinsic excitability. In biological networks these mechanisms are directed to regulate firing activity, and we will adhere to this principle here.

4.3 Functional effects of homeostatic plastic mechanisms

How do the chosen kinds of homeostatic plastic mechanism affect a neuron? Figure 4.1 shows the way in which plasticity of intrinsic excitability and synaptic scaling should affect a hypothetical neuron. We describe the way in which each type of mechanism might regulate firing rate below.

4.3.1 Synaptic scaling

If the overall level of firing of a neuron falls too low, an increase in the strength of all its excitatory connections should help to raise it; similarly, if the level of firing gets too high a decrease in the strength of excitatory connections should help to reduce it. This synaptic change is thought to act with equal force on all afferent synapses to a neuron, i.e., it is a multiplicative scaling of all synaptic strengths [167]. Contrast this with the synapse-specific changes associated with Hebbian correlation-based plastic mechanisms such as long-term potentiation or depression.

The functional effect of synaptic scaling is different when the scaling is applied to the self-connection or to the input connections to a neuron. Effectively, synaptic scaling of the self-connection alters the gain of a neuron, that is, it changes the size of the output response to a given input signal (see Figure 4.1(a)). Synaptic scaling of the input weights affects the influence of the input signal received, and this is reflected by a translation of the output curve along the input axis (see Figure 4.1(b)).

4.3.2 Plasticity of intrinsic excitability

Another way of regulating activity around a constant level is alteration of the intrinsic excitability of neurons. If firing rate is consistently too high, a reduction in the excitability of the neuron should reduce it; if firing rate is consistently too low, increased excitability should raise it. Intrinsic plasticity of the excitability of neurons affects the neuron's function by changing the firing rate response to a given input; effectively this translates the neuron's input/output curve along the input axis (see Figure 4.1(b)).

4.4 Mathematical formulation of homeostatic plastic mechanisms

Homeostatic plasticity is here incorporated into CTRNNs by defining a target range for the firing rates of neurons, corresponding to the postulated set level of activity about which homeostasis is maintained, and then triggering plasticity whenever the firing rate of a neuron is too high or too low. The notion of plasticity activated by high or low firing rates is captured by the use of a plastic facilitation function that varies with firing rate [46]. Two mechanisms were developed to implement homeostatic plasticity in CTRNNs. These were synaptic scaling (Figure 4.1(a)) and an adaptive bias term (Figure 4.1(b)). It should be noted that in biological systems plastic change occurs on a much slower timescale to that of neural firing activity, and that the values of τ_w and τ_b used in Equations (4.2) and



Figure 4.2: Plasticity occurs when firing rate is outside the target range; the size and direction of the excursion determine the rate and direction of plastic change. *Left*: Sigmoidal transfer function showing upper (H_U) and lower (H_L) bounds of target range. *Right*: Plastic facilitation ρ as a function of firing rate z .

(4.3) are chosen to reflect this. Actual values used will be given below where appropriate.

4.4.1 Plastic facilitation

The idea of plastic facilitation was first suggested by Di Paolo in [46]; its use here owes much to his original work, but the form of the function used is slightly different. Plastic facilitation ρ is zero when the firing rate is within the target range and rises or falls linearly to ± 1 outside this range, as defined by Equation 4.1 and shown in Figure 4.2(b).

$$\rho = \begin{cases} \frac{H_L - z}{H_L} & : & 0 \leq z < H_L \\ 0 & : & H_L \leq z \leq H_U \\ \frac{H_U - z}{1 - H_U} & : & H_U < z \leq 1 \end{cases} \quad (4.1)$$

where ρ is the level of plastic facilitation and H_L and H_U are the lower and upper bounds of the target range. It is not clear what the optimal values for these variables would be for any given situation. For most of the experiments reported below, $H_U = 0.75$ and $H_L = 0.25$, but sensitivity tests on these parameters showed that other values could have been used without changing the qualitative nature of the results achieved.

The plastic facilitation function acts in conjunction with the plastic mechanisms described below to give directed plastic change that should regulate firing rate to the target homeostatic range.

4.4.2 Synaptic scaling

When the firing rate of a neuron goes outside the prescribed range all afferent synapses to that neuron (including the self-connection) are multiplicatively scaled. The scaling is directional; it acts so that weights are changed in the direction most likely to bring the neuron firing rate back into bounds. Scaling is applied to both inhibitory (negative weight)

and excitatory (positive weight) synapses. If the firing rate is too high, then excitatory synapse strengths are scaled down and inhibitory synapse strengths are scaled up. If the firing rate is too low, excitatory inputs are scaled up and inhibitory inputs are scaled down. Scaling up or down here refers to the absolute value of the synaptic weight, so that scaling down a negative weight makes it less negative. The size of the change is determined by the plastic facilitation ρ , by a time constant τ_w , and by the current magnitude of the weight. The plasticity rule for synaptic scaling is therefore expressed by Equation 4.2.

$$\tau_w \dot{w} = \rho |w| \quad (4.2)$$

4.4.3 Adaptive bias

Plasticity of the intrinsic excitability of neurons can be implemented in CTRNNs as an adaptive bias term. When a neuron's firing rate goes outside the prescribed range, the bias term of the neuron is shifted to make the neuron more or less likely to fire depending on what is required to bring the firing rate back into bounds. If firing rate is too low the bias is increased, effectively translating the sigmoid activation function so that the neuron is more excitable and hence more likely to fire. If firing rate is too high the bias is decreased, making the neuron less excitable. The size of the change depends on the plastic facilitation ρ and a time constant τ_b . The plasticity rule for intrinsic plasticity is therefore given by Equation 4.3.

$$\tau_b \dot{b} = \rho \quad (4.3)$$

The adaptive bias term used here looks similar to that used in [47] to give neurons that are fast homeostatic oscillators. The main functional differences between the two mechanisms are that here there is a greater separation of timescale between the adapting bias term and the firing rate dynamics (typically at least an order of magnitude difference in time constants compared to same-order time constants in [47]), and also that here the bias adapts only to move firing towards the target homeostatic range, rather than strictly towards a rate of 0.5 as in [47]. If the timescales were the same and the target range was contracted to a single-point range at rate 0.5 then the two mechanisms would be equivalent.

4.5 Can homeostatic plasticity be applied to other kinds of network?

Homeostatic plasticity in biological nervous systems causes changes in neuronal properties depending on the level of activity. For this reason, homeostatic plasticity is most naturally applied to artificial neural networks which incorporate some conception of continuous firing rate and activation dynamics. CTRNNs include these by definition, so are a suitable substrate for the inclusion of homeostatic plasticity. It is possible that homeostatic plastic mechanisms could be adapted for use with other network types (e.g., feed-forward architectures) provided that a target homeostatic range for node output can be defined. For instance, nodes with binary output would not be suitable unless the mean level of activation over some time frame was used as the target for homeostasis.

4.6 Neuroscientific relevance

The plastic mechanisms defined above are obviously highly simplified compared to their biological counterparts, and the use of a plastic facilitation function is a mathematical convenience that cannot be claimed to represent any real biological entity. This means that the relevance of work done with these mechanisms for neuroscience is limited. However, since we are already working with highly abstracted systems (CTRNNs), it is reasonable to hope that no *further* relevance has been lost. In as far as CTRNNs are good models of biological nervous systems, the mechanisms described above are good models of biological homeostatic plasticity. We may not be able to learn much about specific biological mechanisms, but we may be able to understand some of the general qualitative effects of homeostatic plasticity in neural systems. As such, we hope that the mathematical formulation of the plastic mechanisms that is outlined above should be sufficient to let us draw some limited inferences from the work presented in subsequent chapters for neuroscience and for cognitive science.

Chapter 5

Does homeostatic plasticity prevent node saturation?

5.1 Overview

This chapter considers the effects of the homeostatic plastic mechanisms defined in Chapter 4 on the behaviour of a single CTRNN node. It would appear that the homeostatic mechanisms should always stop node saturation from occurring, and in most cases this is true, but the mathematical forms of the CTRNN and of the plastic mechanisms mean that there are cases where the situation is less clear.

5.2 Nature of homeostatic constraint

First of all we will look at the nature of the constraint placed on a CTRNN by the inclusion of homeostatic plasticity, ignoring the plastic mechanisms for a moment and looking just at the significance of the homeostatic target range.

Consider the state equation (Equation 3.1) and activation function (Equation 3.2) for neurons in an N -node network. To meet the homeostatic constraint on activation, all z_i must fall into the target range. The range is defined to be symmetric about $z = 0.5$, and cannot exceed the range for z (recall that $z \in [0, 1]$). So, for some $\delta \in [0, 0.5]$ we have:

$$z \in [0.5 - \delta, 0.5 + \delta]$$

This gives two boundary conditions for the satisfaction of the constraint:

$$z \geq 0.5 - \delta$$

$$z \leq 0.5 + \delta$$

which by substitution into Equation 3.2 and rearranging can be reduced to:

$$-b - \ln\left(\frac{0.5 + \delta}{0.5 - \delta}\right) \leq y \leq -b + \ln\left(\frac{0.5 + \delta}{0.5 - \delta}\right)$$

This must be satisfied at any fixed-point equilibrium state of the homeostatic plastic CTRNN system. For the single-neuron case, the two extremes of the range where this is possible are where the neuron output is equal to the limits of the homeostatic range, i.e., $z = 0.5 - \delta$ and $z = 0.5 + \delta$. At equilibrium, $\dot{y} = 0$. Substituting these values into the state equation (3.1) for a single-neuron network and assuming that there is no external input, we get:

$$b_{min} = -(0.5 - \delta)w - \ln\left(\frac{0.5 + \delta}{0.5 - \delta}\right) \quad (5.1)$$

$$b_{max} = -(0.5 + \delta)w + \ln\left(\frac{0.5 + \delta}{0.5 - \delta}\right) \quad (5.2)$$

as boundary conditions restricting the relation between the self-connection weight w and bias b of the node at a fixed-point equilibrium satisfying the homeostatic constraint for some given δ . Equations 5.1 & 5.2 define two straight lines. These are plotted in Figure 5.1 with $\delta = 0.25$. The two lines define a region of parameter space in which the homeostatic constraint will be satisfied at equilibrium.

For more than one node, the boundary lines are replaced by boundary hyper-planes and the constraint-satisfying region is impossible to represent on the page. However, we hope that the reader will be able to generalise from the one-node example given; the idea of a constraint-satisfying region of parameter space generalises to networks of any size.

For networks of all sizes, the size of the satisficing range is determined by the size of δ (an example value of 0.25 is used in Figure 5.1). The case where $\delta = 0$ is the condition for a centre-crossing network [127] and is plotted in Figure 5.1 for reference.

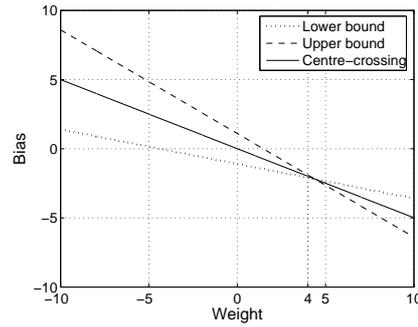


Figure 5.1: Boundaries of region of single-node CTRNN parameter space that will satisfy homeostatic constraint at fixed-point equilibrium. Centre-crossing line also shown. Constraint can be satisfied by fixed-point equilibria when (w, b) lies between lower and upper boundary lines. Note the impossibility of a constraint-satisfying fixed-point equilibrium for $w > 4$.

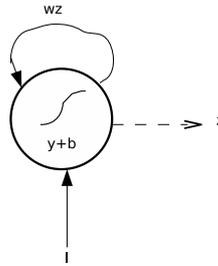


Figure 5.2: Schematic of a single-node CTRNN. The node receives input from itself (wz) and also external input (I). These quantities are summed and contribute to node activation (y), which changes according to Equation 5.3. Node firing rate (z) is a sigmoidal function of activation (y) and bias (b), as specified by Equation 3.2.

5.3 Satisfaction of homeostatic constraint in a single node

In the single-node case (see Figure 5.2), the state equation (Equation 3.1) simplifies to Equation 5.3 below:

$$\tau_y \dot{y} = -y + wz + I \quad (5.3)$$

while the transfer function (Equation 3.2) is unchanged. The simplicity of the single-node case makes it a good example for demonstrating the constraints placed on a CTRNN node by homeostatic plasticity. To do so, here we apply plasticity to an ensemble of nodes and see how the plasticity moves these nodes in parameter space. Because the homeostatic plastic mechanisms we have defined are directed, they always move the nodes towards a constraint-satisfying parameter set if one exists.

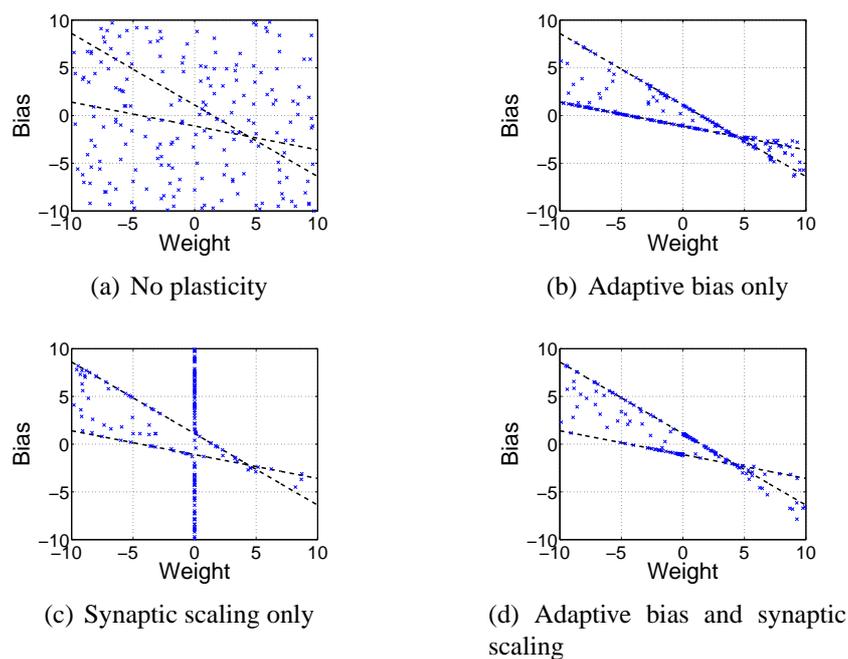


Figure 5.3: Final node positions in parameter space after homeostatic plasticity was applied (with super-imposed boundary lines for constraint-satisfying region). After plasticity most nodes are located in the constraint-satisfying region, with some exceptions. Some nodes are trapped against the vertical line of $w = 0$ by the inability of the synaptic scaling mechanism to change the sign of the synaptic weight. Also some nodes are scattered in a noisy distribution in the bottom-right of the plots, due to an inability to satisfy the constraint.

Parameter sets for 200 single nodes (as specified by Equations 3.2, 4.2, 4.3 and 5.3) were randomly generated ($\tau_y \in [1.00, 4.00]$, $w \in [-10.00, 10.00]$, $b \in [-10.00, 10.00]$, $\tau_w = 40$, $\tau_b = 20$) and instantiated as 1-node CTRNNs. The nodes were then updated for 500 timesteps with zero external input ($I = 0$) while homeostatic plasticity was applied in four different regimes: non-plastic control, synaptic scaling alone, adaptive bias alone, synaptic scaling and adaptive bias together. Figure 5.3 shows the final positions of the nodes in parameter space after homeostatic plasticity has been applied, with the boundary lines from Figure 5.1 superimposed for comparison. Where plasticity has been applied, nodes are mostly located in the constraint-satisfying region of parameter space, with some exceptions. When the synaptic scaling mechanism is used some nodes are trapped against the $w = 0$ line. Also, with both of the homeostatic plastic mechanisms there is a region in the bottom-right hand corner of the plots where nodes are scattered in a noisy distribution. These exceptions are explained below.

Before plasticity is applied, nodes are randomly scattered in parameter space, which can be seen by the positions of the nodes where no plasticity was applied (Figure 5.3(a)).

The adaptive bias mechanism moves the nodes through parameter space by altering the bias, i.e., it moves the representative point vertically in Figure 5.3. The synaptic scaling mechanism alters the synaptic weight, i.e., it moves the representative point horizontally in Figure 5.3. Synaptic scaling has a significant restriction in that it cannot alter the sign of the weight. This means that the synaptic scaling mechanism is not always successful in getting the node into the constraint-satisfying region of parameter space; the inability to alter the sign means that the vertical line of $w = 0$ acts as a barrier trapping the representative point. Combining the two mechanisms means that the node can move in all directions through parameter space (see Figure 5.3(d)).

The noisy distribution of nodes in the bottom-right region of the plots in Figure 5.3 is due to the impossibility of a constraint-satisfying fixed-point equilibrium for $w > 4$. This will be examined in the next section.

5.4 Equilibrium points for the single-node system

Methods for finding the equilibrium states of a dynamical system as a parameter is varied can be found in [160]; sometimes this can be done analytically, but in other cases this is not possible and numerical methods may be used. A useful technique is to plot the set of all the equilibria of a system as a particular parameter is varied. In most cases this forms a connected pathway through parameter space that is traversed as the parameter varies, though in some situations multiple equilibria may exist for a given parameter value.

For the homeostatic constraints to be satisfied in the single-node system, the node must reach a stable equilibrium with its firing rate within bounds. The possibility of this is determined by the shape of the plot of all equilibrium firing rates for the single-node dynamical system, examples of which are shown in Figure 5.4.

Figure 5.4 shows how the equilibrium value of the node firing rate z changes as the external input I changes. Recall that an equilibrium state of a system is a state that will persist as long as the system is not perturbed. A stable equilibrium is robust to minor perturbations; the system will return to the equilibrium position after its state is altered by a small amount. An unstable equilibrium is not robust to perturbations; even a small perturbation will cause the system to diverge from its equilibrium position.

In the single-node system the number of equilibria changes at a bifurcation point at $w = 4$. A bifurcation occurs when a change in some parameter causes a qualitative change in the behaviour of the system. Details of how the bifurcation points of a dynamical system are found are given in [160]. In this case the bifurcation at $w = 4$ separates the unistable and bistable regimes. For $w \leq 4$ the system is in the unistable regime, for $w > 4$ it is bistable. In the unistable regime there is a single stable equilibrium, in the bistable

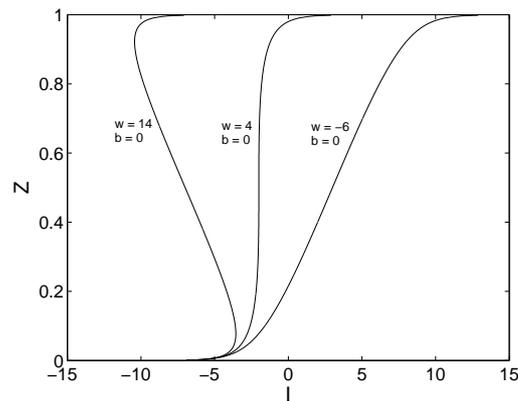


Figure 5.4: Firing rates z at equilibrium for a single node, plotted against external input I for 3 different (w, b) parameter pairs, where w is the self-connection strength and b is the bias term. The plots show how equilibrium firing rate z varies with external input I for $I \in [-15, 15]$. For a single node the value of w determines the slope of the plot; the slope increases as w increases, tilting back on itself for $w > 4$.

regime there are two stable equilibria separated by an unstable equilibrium.

5.5 Behaviour of the single-node system in the unistable and bistable regimes

Figure 5.5 plots equilibrium firing rate against external input for nodes in the unistable and bistable regimes. It might appear that the shape of the plots is determined by the sigmoidal transfer function (Equation 3.2), but the situation is actually a little more complicated. The equilibrium firing rate is determined by the dynamics of the system and its initial conditions, and is not a simple mapping.

While the transfer function obviously limits the equilibrium states to z -values in the range $[0, 1]$, the shape of the plot of all equilibria is actually determined by the strength of the self-connection w . Raising w increases the slope of the plot at its steepest point. The slope of the plot is small for $w < 4$, becomes vertical at $w = 4$, then tilts back on itself to become multi-valued and enter the bistable regime when $w > 4$.

The unistable regime (see Figure 5.5(a)) is quite straightforward to understand. Here z is a single-valued monotonically increasing function of I , w and b . z varies with I , a change in w alters the slope of the plot (and hence the size of response in z to a change in I), while a change in b translates the plot along the horizontal axis while leaving the shape of the curve unchanged. A change in w so that $w > 4$ causes the node to leave the

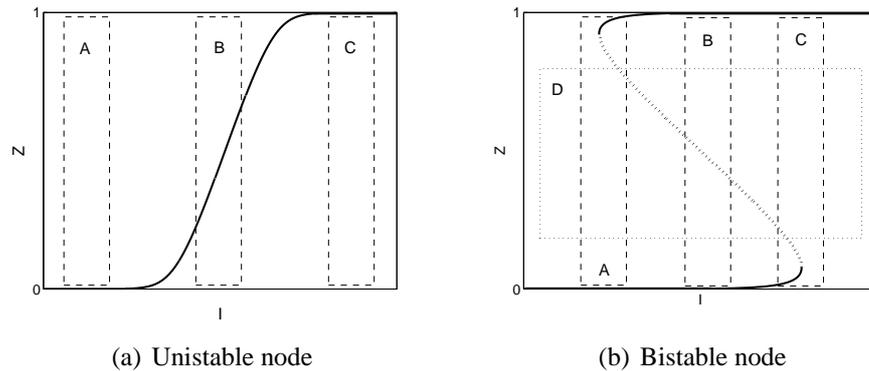


Figure 5.5: Equilibrium firing rates for unistable and bistable nodes showing different input ranges. Plots give equilibrium firing rate z against external input I . Solid line shows a stable equilibrium firing rate, dotted line shows unstable equilibrium firing rate. Shape of the plots is determined by single-node dynamics. Different ranges of input are superimposed on the plots (dashed lines); see text for details of system behaviour when input varies in these ranges.

unistable regime and become bistable.¹

The bistable regime is more complicated. First consider the unstable equilibrium associated with a particular (w, b, I) triple. This equilibrium will only ever be reached if that is where the system is initialised. Because it is unstable, if the system is even a miniscule distance from the equilibrium state it will diverge to one of the stable equilibria.

This is illustrated in Figure 5.6, which plots \dot{y} (the differential of y with respect to time) against y for example nodes in the unistable and bistable regimes. Below each plot is the vector field for the system, which shows what will happen for any value of y . The shape of the curve for the node in the unistable regime (Figure 5.6(a)) is such that it can only intersect the \dot{y} axis at one point. This means that the node has a single stable attractor, so there is only one stable firing rate. The shape of the curve for the node in the bistable regime is such that it may intersect the \dot{y} axis at one, two or three points. This means that the node has one, two or three attractors depending on its parameterisation. The node shown in Figure 5.6(b) has three attractors, one unstable and two stable, so it can fire at two different stable rates.

Changing I translates the plots in Figure 5.6 on the vertical axis without changing the shape of the curve. In the unistable regime this changes the value of the stable firing rate. In the bistable regime it may change the values of the stable firing rates, or may move the system to a point where it has only one stable firing rate.

¹Note that in a multi-node system synaptic scaling of the input weights has a translation effect similar to the adaptive bias mechanism; in this scenario synaptic scaling has both a slope-altering effect caused by scaling of the self-connection and a translation effect caused by scaling of input weights.

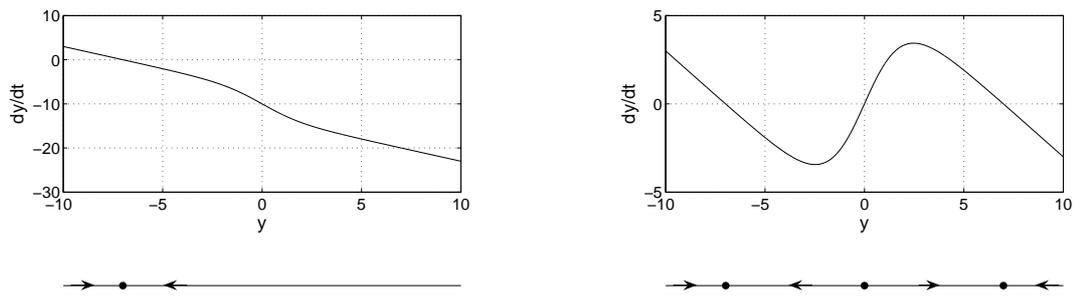
(a) Unistable node $(w, b, I) = (-6, 0, -7)$ (b) Bistable node $(w, b, I) = (14, 0, -7)$

Figure 5.6: Plots of dy/dt against y (top) and vector fields (bottom) for example nodes in the unistable and bistable regimes. In the vector field plots, the black circles on the line represent the attractors (i.e., where $\dot{y} = 0$). The direction of the arrows on the line signifies the sign of \dot{y} ; the system will move to the next attractor in the direction of the arrow. The arrows illustrate the stability of the attractor; if the arrows on either side of the attractor point towards it, it is a stable equilibrium. If the arrows point away from it, it is an unstable equilibrium. There is one stable attractor in the unistable regime, and two stable attractors separated by an unstable attractor in the bistable regime.

Figure 5.5 shows example plots of equilibrium firing rates for unistable and bistable nodes with super-imposed input ranges that provide qualitative examples of the types of behaviour that might be expected when I varies within a range. In Figure 5.5(a), inputs varying in ranges A and C will give rise to a saturated response; range A gives a saturated-low response and range C gives a saturated-high response. Inputs varying in range B will cause a significant change in firing rate, since range B is positioned over the steepest part of the curve.

The situation for the bistable node, shown in Figure 5.5(b) is more complicated. There are six distinct types of system behaviour possible for the bistable node. It is of course possible for the normal input to the node to vary in ranges giving rise to saturated-low and saturated-high responses, similarly to the unistable case, and these input ranges have been omitted from the plot for clarity. These saturated responses are the first two types of behaviour. The next two behaviours exhibit *hysteresis*, that is, the behaviour of the system depends on how it has behaved in the past.

When input varies in range A, the firing rate will be unable to move away from the low-firing equilibrium once it reaches it. It may start at the high-firing equilibrium and then fall to the low-firing equilibrium due to perturbation by the variation in I , but once it reaches the low-firing equilibrium it remains there, with firing rate varying slightly as input changes. A similar situation exists in input range C, but in reverse; the firing rate

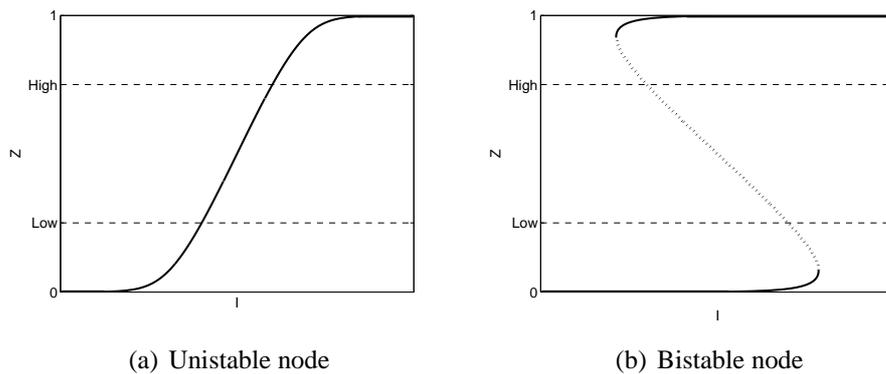


Figure 5.7: Equilibrium firing rates for unistable and bistable nodes showing target range for firing rate homeostasis. Equilibrium firing rate z is plotted against external input I . There is always a stable equilibrium firing rate within the target range when the node is in the unistable regime, but this may not be true in the bistable regime, where it is possible that the only attractors in the target range are unstable and thus unattainable.

may begin low but once it reaches the high-firing equilibrium it remains there. When I varies in the range B the firing rate will remain at whichever equilibrium it starts off at; it is not possible for the firing to flip from high-to-low or low-to-high. The most interesting input range is range D, in which it is possible for the firing to take the low-firing or high-firing equilibria and to flip in between them depending on input.

The significance and utility of node bistability for network behaviour is a topic in itself and will not be examined further here. Suffice to note that bistable nodes may allow an approximation to discrete switching in an otherwise continuous system, if the input to any candidate switching node varies in the region of the transition between the stable equilibria. Such switching may allow a network to alternate between different modes of behaviour depending on input.

5.6 Effect of homeostatic plasticity on equilibria

Of direct relevance to the current study is the observation that there is a clear conflict between homeostatic plasticity driving node activation towards the centre of its range and nodes where the parameterisation does not allow this to happen.

Figure 5.7 shows equilibrium firing rates in relation to the homeostatic range for firing rates. As homeostatic plasticity attempts to drive the firing rate to a stable equilibrium that lies within the target range it alters the relationship between firing rate and external input. The adaptive bias mechanism² translates the equilibrium curve along the I -axis,

²And in the case of a multi-node system, synaptic scaling of the input weights.

while synaptic scaling applied to the self-connection alters its slope. For any given input I_x , homeostatic plasticity attempts to change node parameters so that there is a stable equilibrium inside the target range. This is always possible in the unistable regime since there are always stable equilibria which lie within the range. However, as suggested by Figure 5.7(b), in the bistable regime it is not always possible since there may be no stable equilibria inside the target range. In Figure 5.7(b), the only equilibrium within the target range is unstable (dotted line).

To find a stable fixed-point equilibrium that satisfies the homeostatic constraint, the task for homeostatic plasticity for a node in the unistable regime is to translate and transform the equilibrium curve so that the stable equilibrium position at I_x falls within the homeostatic range. For a node in the bistable regime, homeostatic plasticity must transform node parameters such that it becomes unistable (thus creating a stable equilibrium in the target range) and then translate the equilibria so that this occurs at I_x . If this is not possible then no fixed-point equilibrium will be found.

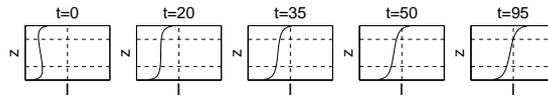
Figures 5.8(a) & 5.9(a) show this process for two examples nodes that start in the bistable regime. Figure 5.8(a) shows homeostatic plasticity (adaptive bias and synaptic scaling mechanisms combined) successfully reaching a fixed-point that satisfies the homeostatic constraint. Figure 5.8(b) shows the bias, weight and firing rate of the same neuron over time. Figures 5.9(a) & 5.9(b) show the same plots for a case where the homeostatic constraint is not met.

What happens in the case where no fixed-point equilibrium can be found is a slow oscillation between low and high firing rates (i.e., a limit cycle), caused by the action of the plastic mechanism. This limit cycle does not satisfy the homeostatic constraint; it is the action of the plasticity continually seeking to do so that causes the oscillation. While in theory it should always be possible to find a fixed-point solution to satisfy the constraint by combining translations and gradient changes of the equilibrium curve, whether or not this actually occurs depends on initial conditions.

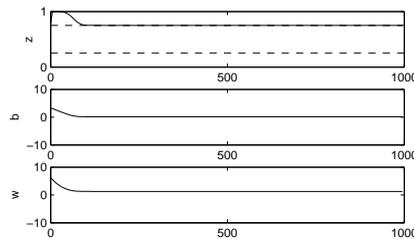
5.7 Constraint satisfaction in a network

So far we have analysed the satisfaction of the homeostatic constraint in the single node case. Regarding networks we have only observed that the idea of a constraint-satisfying region of parameter space is still valid. We will not perform any analysis of constraint-satisfaction in networks, because the mathematics in the N-dimensional case gets rather complicated and does not justify the meagre explanatory pay-off. However, there are a few observations we can make concerning constraint-satisfaction in networks.

The first is to note that each node in the network will independently try and maintain

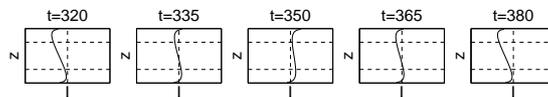


(a) Equilibria at different times. Plots show equilibrium firing rate z against external input I at particular times taken from the time series plots below.

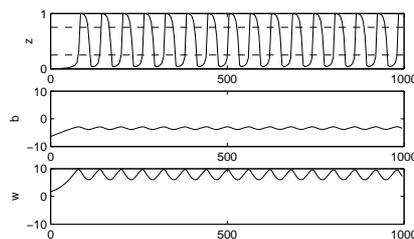


(b) Firing rate z , bias b and self-connection weight w , plotted against time t .

Figure 5.8: Homeostatic plasticity moves a bistable neuron to the unistable regime and finds a constraint-satisfying fixed point solution. Figure 5.8(a) shows equilibria changing over a short period of the time series plots of node parameters and firing rate shown in Figure 5.8(b). Synaptic scaling of the self-connection alters the slope of the equilibrium curve and adaptive bias translates the curve so that it has a stable equilibrium firing rate inside the target range.



(a) Equilibria at different times. Plots show equilibrium firing rate z against external input I at particular times taken from the time series plots below.



(b) Firing rate z , bias b and self-connection weight w , plotted against time t .

Figure 5.9: Homeostatic plasticity fails to make a bistable neuron unistable and finds a limit cycle equilibrium that does not satisfy the homeostatic constraint. Figure 5.9(a) shows equilibria changing over a short period of the time series plots of node parameters and firing rate shown in Figure 5.9(b). Both synaptic scaling and adaptive bias are active, but are unable to find a (w, b) pair where the node has a stable equilibrium firing rate inside the target range. The continuing action of the plasticity results in a stable oscillation where firing rate continually overshoots the target range, with associated oscillatory changes in w and b .

local homeostasis, so that homeostasis in the network is achieved in a decentralised manner. The constraint-satisfaction task is harder in networks than in a single node, because the input each node receives from other nodes. Also, in any functioning network, the continual perturbation of sensory input makes the task harder still. Because of these interactions, the case where a constraint-satisfying equilibrium is found for the whole network is likely to occur infrequently.

In networks, satisfaction of the homeostatic constraint may be better interpreted as a tendency or process, rather than an achievable goal. The parameters of each node in the network are continuously acted on by homeostatic plasticity to move the node towards the constraint-satisfying condition. However, the conditions within which each node operates may change too quickly for this goal ever to be reached.

The same effects that are seen at the level of a single node will also occur in networks. For instance, the slow oscillations caused by the failure to find a fixed-point solution to the homeostatic constraint may be present in any of the nodes in the network. The switching behaviours and slow dynamics this may cause at the network level may be useful or disruptive to a functioning network; this remains a topic for investigation.

5.8 Implications for neuroscience

The analysis above is based entirely on the mathematical formulation of the CTRNNs and homeostatic plastic mechanisms we defined. As such, the results shown have no relevance to biological brains. However, we can observe that the ideas that synaptic scaling of the self-connection alters the slope of the transfer function, and that the adaptive bias (or scaling of input weights) translates the transfer function, may be analogous to the operation of biological neurons; logically something similar must occur. Also, the idea that homeostatic plasticity in neural networks is an ongoing process that is goal-oriented but does not necessarily reach its target is likely to be valid in the biological case.

Chapter 6

What effect does homeostatic plasticity have on network behaviour?

6.1 Overview

Having examined the effect of homeostatic plasticity on a single CTRNN node, now we will consider the effects on a network. In this section we are concerned with typifying the way in which homeostatic plastic CTRNNs are different from non-plastic CTRNNs. To do so we will not consider performance on any particular network function or behavioural task, but rather will attempt to quantify differences in basic network properties: sensitivity, signal propagation and the propensity to display oscillations. In order to fully understand the effects of homeostatic plasticity at a network level, it is necessary also to consider the effects at levels below (node) and above (agent) the network level.

This chapter will first of all present a brief discussion of some methodological issues, followed by the actual methods used and results achieved, before concluding with some discussion of the results.

6.2 Methodology

There are a number of methodological issues that need to be considered to ensure that a reasonable understanding of the effects of homeostatic plasticity is obtained. These range from the problems of fair comparison between different types of neural network to those

of defining an appropriate level of analysis. Some discussion of these issues is given in this section.

6.2.1 Plasticity, learning, and progress

Before proceeding it is worth a brief clarification of terminology. There is a tendency in the literature to equate plasticity with learning, but this adds a subjective impression of progress to the concept that is not necessarily grounded in observed data. *Learning* implies the acquisition of a new skill or goal-directed behaviour, while *plasticity* does not. Learning may involve plasticity, but plasticity does not imply learning.

In this thesis, we use the term ‘plasticity’ to refer to online (or ‘lifetime’) changes in the weights or biases of a neural network. By ‘plastic mechanism’ we mean the mechanism by which such changes are effected. Thus our conception of plasticity implies no sense of improvement and is independent of its utility to a host agent. We realise that our definition makes a distinction between ‘plastic’ variables (i.e., weights and biases) and other variables (e.g., firing rates) that is largely artificial — they are all just variables in a dynamical system — but feel that adopting this terminology allows for a clearer discussion.

6.2.2 Measuring the effects of plasticity

Some examples of previous uses of plasticity in CTRNNs include the application of Hebbian rules to controllers for a robot navigation task [60], the combination of Hebbian rules with Ashbyan homeostatic adaptation [46], and the use of mechanisms analogous to the action of diffusing gases [87] (which can be argued to be a form of plasticity since the gas alters the way in which firing rates interact and thus creates a second order effect).

It is important to remember that adding plasticity to a CTRNN is not a simple modification or extension of it, but makes a significant change to network dynamics. It is tempting to see plasticity as a quantitative alteration of network parameters akin to a mutation in the network genotype, but this view is misleading. Adding plasticity to a CTRNN makes a qualitative change to its dynamics and can create a new class of controller with very different characteristics.

Plasticity turns parameters into variables. For example, Hebbian plasticity rules change the strength of synapses dependent on node activation. This means that in a Hebbian network the synaptic weights are variables, whereas in a non-plastic network they are parameters. Plasticity may act on a slower timescale when compared with firing rates, but the increased number of variables adds extra degrees of freedom to the system irrespective of

timescale. The corresponding increase in the dimensionality of the network's dynamics will have implications for evolvability.

How should we assess the usefulness of different kinds of plasticity? Given the general aim of designing useful robots, comparative studies of the performance of different types of controller in various evolutionary trials are a valid approach (some evolutionary experiments will be reported in Chapter 7). However, the evolutionary approach is limited. Evolutionary trials are by their nature task-dependent and we should not be surprised if one sort of controller outperforms another in a particular trial, but success in one scenario does not necessarily imply superiority overall. Different sorts of controller have different sorts of dynamics, and may be better suited to some tasks than to others; comparison of evolvability is only possible when the performance criteria are well-specified.¹

In pursuit of a more general comparison, it is useful to view plastic neural controllers as dynamical systems composed of a coupling between a firing rate subsystem and a plasticity subsystem. The way in which these subsystems interact and give rise to global dynamics may allow us to classify the intrinsic dynamics of different kinds of neural controller and find qualitative differences that allow us to match controllers with tasks for which they will be well-suited. By this approach we may gain a theoretical base to complement and guide the evolutionary robotics design methodology, moving away from intuition and 'black box' design.

In this chapter, we look at the effect of homeostatic plasticity on some properties of CTRNNs that are not task-specific, but instead might be significant for a variety of network tasks. We hope to gain a general understanding of the effects of homeostatic plasticity that will help to explain the results gained from evolutionary trials.

6.2.3 Adiabatic approximation

Plasticity in neural systems typically occurs on a much slower timescale to that of firing rate dynamics, often by several orders of magnitude. This allows us to make a simplifying assumption and treat the slower-changing plastic variables as fixed in relation to the faster-changing firing rate variables, on the basis that the faster variables will typically reach an equilibrium before the slower variables have changed significantly. This approach is based on the technique of adiabatic elimination [73] and gives useful purchase on the analysis of network behaviour.

We can approximate the behaviour of the whole plastic system by studying the behaviour of non-plastic snapshots of the system at particular instants in time, treating the

¹Goats are better at climbing mountains while sheep provide more wool, but it is hard to say that one species is better than the other.

slow plastic variables as fixed parameters. Effectively we consider the state trajectory of the homeostatic plastic CTRNN as a journey through CTRNN-space. The state of the plastic network at any point in time can be used to specify a non-plastic CTRNN. Comparison of the non-plastic networks instantiated at different points on this trajectory should give some indication of the behaviour of the plastic network. Importantly, adiabatic elimination can allow a more direct comparison of plastic networks to non-plastic networks in some cases, because a non-plastic CTRNN can be instantiated from a plastic CTRNN by freezing the plastic variables at any particular point in time.

The technique of adiabatic elimination is a useful tool and will be used extensively in the work presented in this chapter, but it does not tell the whole story. We have already looked at the dynamics of homeostatic plasticity in a single node in Chapter 5, and it is important also to conduct analysis of the dynamics and behaviour of fully functioning homeostatic plastic networks. However, mathematical analysis of larger networks is problematic, and the method of looking at the change in behaviour of non-plastic CTRNNs instantiated from the attributes of plastic networks as they move through CTRNN-space allows useful comparisons to be made.

6.2.4 Lifetime plasticity versus developmental period

Another issue in assessing the effect of homeostatic plasticity concerns the time at which the plasticity should be applied. Should the plastic mechanism be used as a developmental mechanism, and applied for a lengthy continuous period before the network is assessed? Or should plasticity be applied in an ongoing manner during the entire functional lifetime of the network, with assessment right from the start?

This issue is a greater concern when evolutionary trials are used, and will be treated in more depth in Chapter 7. For the current set of experiments (that seek to determine what effect homeostatic plasticity has on generic network dynamics) there is no network ‘function’ or ‘lifetime’ in any meaningful sense. We are more concerned with network behaviour before, after, and during the application of homeostatic plasticity. Experiments in this section are performed with plasticity applied continuously unless otherwise stated.

6.2.5 Measurement from random ensembles

We want to determine the effects of homeostatic plasticity when applied to CTRNNs in a general sense. Therefore we are not interested in the effects of the plasticity on any particular individual CTRNN or on CTRNNs trained to perform a particular task. We wish to understand how homeostatic plasticity affects CTRNNs in general, which we

hope will then allow us to predict what will happen when homeostatic plasticity is applied to individual networks in particular.

To understand the general effects of homeostatic plasticity on CTRNNs and the behaviour of homeostatic networks, we need to look across the whole space of possible networks. Looking at networks trained for a particular task restricts the study to a particular region of parameter space. Here we wish to expand our focus to cover the entire space. The large number of parameters of CTRNNs means that exhaustive sampling of the space is not feasible. Instead we will look at large ensembles of networks instantiated from randomly chosen points in the parameter space. By looking at large ensembles it should be possible to get a representative picture of the whole space.

6.2.6 Metrics for measuring networks

Because we are not interested in particular functions or behaviours, but instead want to look at generic effects, the metrics we use to quantify the effects of plasticity are necessarily very general. We discussed above the aim of improving CTRNNs as an evolutionary substrate, and with this in mind the metrics we will use are some that we hope will be relevant to the evolvability of neural robot controllers in general, without being specific to any particular problem.

6.2.6.1 Useful qualities for behavioural substrates

Without performing any evolutionary experiments, we might look for types of controller that make it easy to generate behaviour in an autonomous agent placed in an environment. Envisaging basic robotic tasks such as navigating a real-world environment, avoiding obstacles, locating opportunities, and escaping threats, and observing that such behaviours will form the basis of any more advanced tasks a mobile autonomous agent might be asked to carry out, we feel that a good behavioural substrate will be one rich in the sorts of behavioural primitive that might support such actions.

First, the agent must be capable of reacting to its environment. This may sound over-simple, but it is not achieved in saturated CTRNNs. Reacting to the environment requires that the agent be possessed of sensors and actuators with an effective link between them. Signals must be transduced in some form from input to output to complete the sensorimotor loop [28, 30]. In a neural network controller, the input nodes must communicate with the output nodes in some fashion. If they do not, then there is no link between sensory information and action and thus the agent cannot react to a changing environment in any meaningful way. Any action taken would be based solely on internal state with no reference to the outside world. This type of behaviour is unlikely to be adaptive, since it will

not allow opportunities to be exploited nor threats avoided.

Second, and more controversially, for an agent to display interesting or ‘cognitive’ behaviour it must have some form of internal dynamics. Without internal dynamics the agent would be restricted to actions that are purely reactive, that is, it can only respond to instantaneous stimuli and cannot integrate stimuli over time. To do so requires that some internal state variable can store or reflect stimuli received at a previous moment in time. If an agent controller has an ongoing internal dynamic, the task of behaviour generation becomes one of modulating activity and guiding the dynamics to suit the current situation rather than that of generating behaviour from scratch at each moment in time. Sensory input is then seen as a perturbation to an existing process,² rather than as the sole trigger for action.³

Third, building on the previous property, it is widely thought that autonomous oscillations are important in the generation of behaviour. A huge number of biological systems depend on oscillations in some form, from legged locomotion to digestion, and it is believed that many of these phenomena depend on rhythmic patterns generated in the nervous system. Central pattern generation has thus been the focus of much research effort, and it is not unreasonable to suppose that the ability to produce autonomous oscillations will be important to artificial agents as well as biological ones. [67] offers a good discussion of physiological rhythms, while [88], [29] and [17] describe good examples of this kind of research from the field of adaptive behaviour.

6.2.6.2 Metrics

The properties described in the previous section form the basis of our metrics for quantifying network behaviour. We will measure the sensitivity of nodes in the network, how easily signals are propagated, and how likely networks are to display oscillatory dynamics. By measuring the performance of ensembles of different sorts of network on these metrics we will be able to make qualitative distinctions between plastic and non-plastic networks and understand the effect of homeostatic plasticity on CTRNNs in general.

6.2.7 Levels of analysis

When seeking to understand the effect of plasticity acting on a neural network, it is useful to look at behaviour at the levels immediately below and above the network, that is, at

²Cf. dynamical systems approaches to cognition and behaviour [14, 16, 164, 173]

³A counter-example to this argument is provided by stigmergic agents such as social insects, whose behaviour is guided entirely by their environment. However, it may be argued that stigmergy involves a form of memory just as much as other kinds of behaviour that are more commonly described as ‘cognitive’; it is just that in the case of stigmergic behaviour the memory is *external* to the agent.

the levels of the node and the agent. By understanding what occurs in individual nodes we can make predictions about the network level and better explain what is observed. By considering the level of the agent that is controlled by a network, we can interpret the network level properties we see in terms of the actions of an agent embedded in an environment. In the experiments reported below we consider the effects of plasticity at the levels of node, network and embedded agent.

6.3 Method

In the set of experiments reported below the neural network and homeostatic plasticity are implemented as described in Section 3.4.1 and Section 4.4, and this method will not be repeated here. In this section we describe a robot simulation platform used to look at agent-level behaviours and the implementation of the metrics used to measure network performance.

6.3.1 Photo-sensitive robot simulator

To give an agent-level illustration of the different behaviours generated by plastic and non-plastic control networks, a simple simulation of a light-sensitive robot was developed. The simulated robot is not intended to represent any real-world robot, so there can be no intention of transferring controllers from simulation to reality and we are able to sidestep many of the problems that this can cause.

The robot is modelled as a circular body with two motors mounted at either ends of an axle along its diameter and two light sensors mounted at angles of $\pm \frac{\pi}{3}$ radians from the forward direction, as shown in Figure 6.1.

6.3.1.1 Kinematics

The simulated robot kinematics are adapted from [56]. Differential drive steering is accomplished by the two motors, which may give thrust both forwards and backwards, allowing the robot to spin while stationary or to move forwards or backwards with any instantaneous angular velocity. Robots are assumed to have negligible mass, so that the motor output can be taken as the tangential velocity of the robot at the motor mount point. The instantaneous angular velocity ω of the robot is calculated using Equation 6.1, and this is used to calculate the motion of the robot using Equation 6.2, where V_L and V_R are the instantaneous velocities of the left and right motors respectively, d is the length of the axle, ICC denotes the instantaneous centre of curvature (the imaginary point about which the robot orbits at any given instant), β is the current heading of the robot, and (x, y) is the

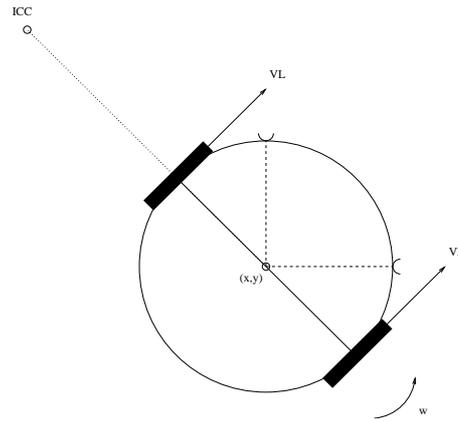


Figure 6.1: Diagram of photosensitive robot. The robot has two wheels that give differential drive steering in response to sensory input from its two sensors.

robot's position. Details of how this system is derived can be found in [56]. The equations are updated using Euler's forward method with a timestep equal to that used for the neural controller.

$$\omega = (V_R - V_L)/d \quad (6.1)$$

$$\begin{aligned} x' &= (x - ICC_x)\cos(\omega\delta t) - (y - ICC_y)\sin(\omega\delta t) + ICC_x \\ y' &= (x - ICC_x)\sin(\omega\delta t) + (y - ICC_y)\cos(\omega\delta t) + ICC_y \\ \beta' &= \beta + \omega\delta t \end{aligned} \quad (6.2)$$

6.3.1.2 Light sensors

The light sensors return a signal with strength inversely proportional to the distance D of the sensor from the light source. This is modelled as a signal varying as an inverse exponential function (Equation 6.3). Sensors are assumed to be mounted on top of the agent, so that they are never in the shadow of the robot's body, and have a maximum range at which they can detect light (denoted by D_{max} and set to 500 for the experiments reported in this chapter).

$$Signal = e^{\frac{-D}{D_{max}}} \quad (6.3)$$

6.3.1.3 Control network

The robot is controlled by a fully connected 6-node CTRNN of the type described above. Two nodes are designated sensor nodes and each receive the signal from one of the light sensors modified by an agent-symmetrical gain parameter (drawn from the range $[0, 10.00]$) as external input. Two other nodes are designated motor nodes, and the output from these nodes is scaled to the range $[-1.00, 1.00]$ and modified by an agent-symmetrical gain (again drawn from the range $[0, 10.00]$) to give the instantaneous velocity created by that motor. Both sensory input and motor output have noise added before the gain is applied, drawn from a uniform distribution in the range $[-0.25, 0.25]$.

6.3.1.4 Environment

Robots are placed into an infinite featureless plane containing one or more point light sources. The sensory input from these sources is determined by Equation 6.3, as mentioned previously.

6.3.2 Metrics

The metrics used to measure network performance are how well signals propagate through the network, and how likely the network is to display oscillatory dynamics. These metrics are supported by measurement of the sensitivity of individual nodes, and by some observation of the behaviours of photo-sensitive agents. The metrics will be used to measure the effect of homeostatic plasticity in and on CTRNNs.

6.3.2.1 Node sensitivity

We here define sensitivity as the change in output caused by a change in input. More specifically, we look at the mean change in output caused by a series of random step changes in the external input to the node. Sensitivity is calculated here on nodes with a self-connection, i.e., on the type of nodes used in fully connected CTRNNs.

To test sensitivity, external input I (randomly drawn from a uniform distribution in the range $[-1.00, 1.00]$) is applied to the node and held constant for a period (usually 20 timesteps) to allow the node to reach a stable equilibrium output, at which point a new external input is randomly selected from the same range. The firing rate of the node at equilibrium is recorded immediately before each change in input is applied and used to calculate the absolute difference in firing rate resulting from each change in input. Sensitivity is defined as the mean change in firing rate caused by 1000 successive random input changes, as given in Equation 6.4.

$$Sensitivity = \frac{1}{1000} \sum_{n=1}^{1000} |z_n - z_{n-1}| \quad (6.4)$$

6.3.2.2 Signal propagation

Signal propagation is calculated similarly to node sensitivity, but on network architectures. We want to get a picture of how homeostatic plasticity affects signal propagation in all types of network architecture, but exhaustive coverage of the whole space of possible topologies lies beyond the scope of this paper. Instead we use a subset of topologies (shown in Figure 6.2) that we feel is sufficient to gain a general understanding. To measure the impact of homeostatic plasticity, we will look at signal propagation in networks before and after a period of homeostatic plasticity; this period of plasticity might be loosely compared with a developmental phase in a biological organism. Data will always be gathered on non-plastic networks in order to avoid any interference from the plastic mechanisms and allow a fair comparison.

Signal propagation will be measured as the mean change in firing rates in a network caused by a change in input. In order to gain a representative measure of signal propagation in a particular type of CTRNN topology, we will look at large numbers of randomly parameterised networks of that type, measuring the mean change in firing caused by many different input changes.

This conception of signal propagation is rudimentary and does not consider the information content of a signal or transmission accuracy. Nor do we consider the prospect of an appropriately adapted CTRNN performing computation over its inputs such that in some cases the output would not change for a given change in sensory input (e.g., a predator that has evolved to do nothing when presented with inappropriate prey). However, the point we are trying to make concerns the very possibility of *any* information being transmitted, and as such we feel that such a basic measure is justified.

6.3.2.3 Oscillations

The mathematical formulation of CTRNNs means that oscillatory dynamics are not possible in a single node, but are only possible in networks of 2 or more nodes depending on parameterisation. In a large network there are many subcircuits that could maintain a stable oscillation; in an N-node network there are $\frac{N(N-1)}{2}$ 2-node subcircuits, $\frac{N(N-1)(N-2)}{6}$ 3-node subcircuits, and so on. As N rises, there is thus an exponential increase in the number of possible oscillatory subcircuits and we would expect the likelihood of an au-

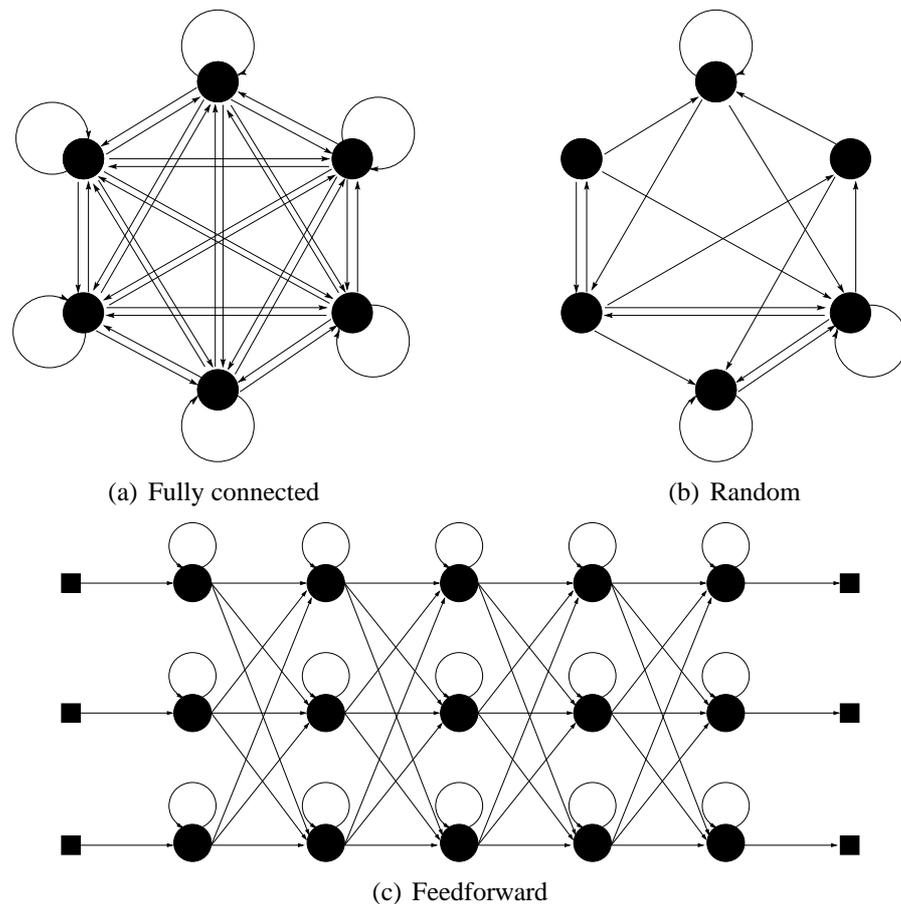


Figure 6.2: Different network topologies used to test signal propagation: (a) Fully connected — each node has a self-connection and an afferent connection from every other node, (b) Random — based on Erdős-Renyi scheme [57], example shown has 50% probability of connection between each pair of nodes, (c) Feedforward — there are no backwards or lateral connections. External input is applied to an arbitrarily chosen subset of nodes in the fully connected and randomly connected network topologies, or to the first layer in the feedforward topology.

tonomous oscillatory dynamic to rise.⁴

Whether or not an oscillatory stable state does in fact occur is dependent on network parameters (connection weights, biases, decay constants) and on the initial state of the system. Some network parameterisations will always produce an oscillation irrespective of initialisation, some will never produce an oscillation, and many will show oscillations if the initial condition is favourable.

We looked at how likely CTRNNs were to display autonomous oscillations. Specifically, we considered the statistical likelihood of an oscillatory dynamic occurring in a randomly parameterised network. This was measured by creating large ensembles of random CTRNNs, starting each network in a variety of different initial conditions, and observing its subsequent behaviour to see if an oscillation occurred in any subcircuit of the network.

Oscillations were detected by looking at a combination of variance in neuron firing rate and repeated change of sign in the rate of change of neuron potential, since these together indicate an oscillatory dynamic; repeated changes of sign of the rate of change of potential mean that there is an oscillation in potential, while non-zero variance in firing rate means that this oscillation in potential is translated into an oscillation in firing rate (it is possible that a saturated node may display an oscillation in potential without this being detected in its firing rate). This method is crude and does not differentiate between limit cycle and chaotic oscillations, but is computationally tractable and thus allows sufficient data to be collected in a reasonable time.

6.4 Node-level effects of homeostatic plasticity

6.4.1 Slow oscillations

The first node-level effect to observe is that the introduction of homeostatic plasticity creates the possibility of slow oscillations in firing rate as the plastic mechanism causes firing rate to repeatedly flip-flop between too-high and too-low levels. This effect is contingent on initialisation and is a result of the lack of a stable firing rate equilibrium inside the target homeostatic range. This effect was discussed earlier (in Section 5.6) and will not be further described here.

⁴The precise nature of the relationship between network size and the likelihood of oscillatory dynamics is not straightforward to determine and will be affected by connectivity [95].

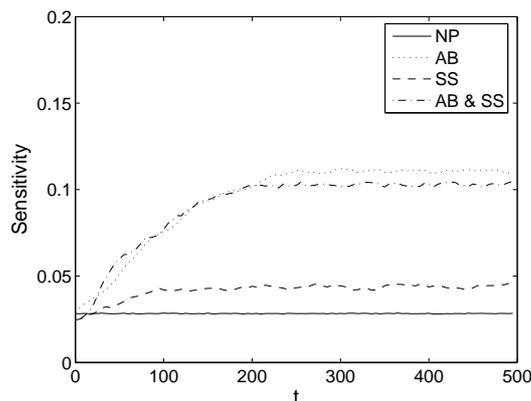


Figure 6.3: Node sensitivity rises over time as homeostatic plasticity is applied. NP - no plasticity, AB - adaptive bias, SS - synaptic scaling.

6.4.2 Increased sensitivity

Parameter sets for 200 single nodes were randomly generated ($\tau_y \in [1.00, 4.00]$, $w \in [-10.00, 10.00]$, $b \in [-10.00, 10.00]$, $\tau_w = 40$, $\tau_b = 20$) and instantiated. The nodes were then updated for 500 timesteps while homeostatic plasticity was applied with different regimes (synaptic scaling alone, adaptive bias alone, both synaptic scaling and adaptive bias together). Every 5 timesteps a snapshot of node parameters was taken and instantiated in a non-plastic node in order to measure changes in node sensitivity over time as plasticity was applied.

Sensitivity is plotted against time in figure 6.3, which shows that the homeostatic plastic mechanisms cause a rise in node sensitivity over time. The adaptive bias mechanism is significantly more effective in doing so than the synaptic scaling mechanism, while both mechanisms together are the most effective.

The reason for the greater efficacy of the adaptive bias mechanism over the synaptic scaling mechanism is hinted at in Figure 5.3. Recall that homeostatic plastic change continues until the representative point of the node in parameter space lies within the constraint-satisfying region. Also recall that the adaptive bias mechanism always successfully achieved a constraint-satisfying parameter set, whereas the synaptic scaling mechanism cannot alter the sign of the weight and thus is not always successful in satisfying the constraints. Since sensitivity is maximised when nodes are in the constraint-satisfying region, this explains the low efficacy of the synaptic scaling mechanism in raising node sensitivity compared to the adaptive bias mechanism. Combining the two mechanisms means that the node can move in all directions through parameter space and gives similar sensitivity to the adaptive bias mechanism.

6.5 Network-level effects of homeostatic plasticity

6.5.1 Additional behavioural timescales

The slower timescale introduced by the plasticity, which can cause slow oscillations at the node level, also adds an additional slower timescale to the dynamics at the network level. The slow node-level oscillations still occur when nodes are connected in a network, and can cause bifurcations in network behaviour. For example, a node showing the slow flip-flop behaviour may trigger a switch between a stable oscillatory dynamic in the network when its firing output is high and a fixed-point network attractor when its output is low.

6.5.2 Transient dynamics

As mentioned previously, adding homeostatic plasticity to a CTRNN converts weight and bias parameters into variables, albeit with a slower timescale than firing activity. This increase in the dimensionality of the system has a big effect on the complexity of the dynamics. One way in which this is expressed is in the time taken for the system to converge to an attractor when started from a random initial state. A rough measure of the time taken to converge to an attractor is given by the time taken for the system to converge to a repeating sequence. A repeating sequence in a deterministic system indicates either a fixed point or a limit cycle attractor, and while this measure would not detect convergence to a chaotic attractor (which would not repeat), it serves as a useful metric for comparing the time to convergence of networks with different plasticity regimes. Tables 6.1, 6.2 and 6.3 show results from a series of experiments in which we generated ensembles of randomly parameterised, randomly initialised networks and recorded how long they took to converge to a repeating sequence.

We tested non-plastic CTRNNs and HP-CTRNNs with different sizes of target range for the firing rate and different values for τ_b and τ_w . Ensembles of 1000 networks of each type were randomly initialised with $w, b \in [-10.00, 10.00]$, $\tau_y \in [1.00, 5.00]$ and $y \in [-8.00, 8.00]$, and were updated using Euler's forward method. We monitored network state as it changed over time and measured the lengths of the transient trajectories prior to a repeating sequence being reached, terminating the run after 20000 timesteps if no repetition of state had been found. The results shown include the mean and median lengths of the transients, calculated on those cases when a repeating sequence was reached.

i

The results shown indicate that there is a huge difference in the time to convergence of the different network types. Non-plastic CTRNNs always converged in the given time period, typically taking less than 100 timesteps to do so. In contrast, HP-CTRNNs were

Repeating seq. found	Mean trans. length	Median trans. length
1000/1000	63	47

Table 6.1: Convergence to repeating sequence: Non-plastic 10-node CTRNNs. Transient lengths shown to nearest timestep. Non-plastic CTRNNs converge reliably and quickly.

	Repeating seq. found	Mean trans. length	Median trans. length
$\tau_w = 4, \tau_b = 2$	953/1000	1671	1155
$\tau_w = 8, \tau_b = 4$	919/1000	3513	2670
$\tau_w = 40, \tau_b = 20$	186/1000	10135	11451

Table 6.2: Convergence to repeating sequence: 10-node HP-CTRNNs with target range $z \in [0.05, 0.95]$. Transient lengths shown to nearest timestep. Not all of the plastic networks converged to a repeating sequence in less than the 20000 timesteps available in the trial, with the rate of plastic change having an impact on the likelihood of a repeating sequence being reached. Slower plasticity (larger time constants) made networks less likely to converge and increased the length of the pre-convergence transient period.

	Repeating seq. found	Mean trans. length	Median trans. length
$\tau_w = 4, \tau_b = 2$	637/1000	7444	6327
$\tau_w = 8, \tau_b = 4$	492/1000	10114	9547
$\tau_w = 40, \tau_b = 20$	1/1000	—	—

Table 6.3: Convergence to repeating sequence: 10-node HP-CTRNNs with target range $z \in [0.25, 0.75]$. Transient lengths shown to nearest timestep. These networks were tested with a tighter homeostatic target range than the networks described in Table 6.2, which accentuated the effects of the plasticity on the convergence of these networks. The tighter range made the networks significantly less likely to reach a repeating sequence in the trial duration of 20000 timesteps, and where a repeating sequence was reached the transient period was substantially longer. When plasticity was used at the slowest rate shown ($\tau_w = 40, \tau_b = 20$), only one network reached a repeating sequence in the given time; mean and median transient lengths were not computed for this case due to the paucity of data. It is possible that given more time more networks would have converged, but here the time allowed was restricted to allow fair comparison with data from Tables 6.1 and 6.2.

less likely to reach a repeating state in less than 20000 timesteps, and took many thousands of timesteps if they did so. The tightness of the homeostatic target range had an effect on convergence, as did the rate at which the plasticity acted, with faster plasticity and a more relaxed range giving quicker convergence.

While the plastic networks take a long time to converge to a repeating sequence, it should be noted that the character of their dynamics during the transient phase is not without structure. After a noisy initial stage, they typically show a slowly attenuating convergence that may be best described as ‘spiralling in’ towards the attractor state.

6.5.3 Improved signal propagation

The increased sensitivity caused by homeostatic plasticity at the node level is also seen at the network level, where it is exhibited as an increase in the strength of the signal propagated when network input changes. The effect of homeostatic plasticity on signal propagation was measured for the different types of network topology mentioned above.

Each type of network topology was tested in a similar way. Ensembles of randomly parameterised networks were generated. Signal propagation was examined in each network in its original state, and then again after a period in which homeostatic plasticity was applied. All connection weights and biases were drawn from the range $[-10.00, 10.00]$ and all neuron potential decay constants were drawn from the range $[1.00, 4.00]$. Plasticity was applied with $\tau_w = 40$ and $\tau_b = 20$. Networks were updated using Euler integration with a step size of 0.2 timesteps.

Network input was chosen randomly from a uniform distribution. Input was held constant for a period during which network firing rates were measured and then a new input was chosen. The mean size of the change in output caused by each change in input was calculated over a large number of input presentations to give a representative measure of the change in network activity that might typically be expected from a change in input. This measure was used to compare signal propagation between different network topologies before and after plasticity was applied.

The change in output was assessed in different (though similar) ways depending on the network topology. The mathematical formulation of CTRNNs means that stable oscillatory dynamics can occur if two or more nodes are connected in a loop. In networks where cyclic paths are possible (that is, fully and randomly connected networks but not feed-forward networks), there is a good chance that oscillatory dynamics will occur. In these networks, the mean firing rate was measured for each node over the full period for which input was held constant. The period was chosen to be long enough such that any transient dynamics (while the network settled to a new stable state following the change in input) would have an insignificant effect on the recorded mean. The use of the mean allowed a rough comparison between oscillatory and fixed-point stable states; if the oscillation changed then its mean value would most likely change also. For the feedforward networks, where oscillations cannot occur, the firing rate of each node was measured at the end of each period prior to the presentation of a new input vector. This allowed the network to settle to a new fixed-point before measurement.

Results given here are from simulations where both kinds of homeostatic plasticity (synaptic scaling and adaptive bias) were applied simultaneously. Equivalent runs performed with each mechanism acting alone gave qualitatively similar results. Data from

these runs are omitted for clarity.

6.5.3.1 Fully connected networks

Fully connected networks have a connection in both directions between each pair of nodes, and every node also has a self-connection (see Figure 6.2(a)). They are the type of neural architecture most commonly used in the evolutionary robotics literature. A key point to note is that there is a path length of one link between any pair of nodes, meaning that in a fully connected architecture there is a direct connection between input and output nodes. This direct influence of the input node on the output node is modulated by the activity of the other nodes.

Networks were created with 1, 3, 5 and 10 nodes; 200, 600, 1000 and 2000 networks of each respective size were created to reflect the combinatorial expansion in the number of parameters.⁵ A single node in each network was designated the input node and received input randomly drawn from the range $[-5.00, 5.00]$, held fixed for 200 timesteps. Signal propagation was measured over 1000 input presentations. Homeostatic plasticity was applied for 500 timesteps and then signal propagation was measured again.

Results are plotted in Figure 6.4, which shows the mean changes in output caused by a change in network input for each node in the network, before and after the homeostatic plasticity is applied. The input node in each network is marked on the plots; the other nodes shown are the hidden nodes. The input node shows a significantly greater response than the other nodes. However, before plasticity is applied, even the input node does not demonstrate a large mean change in state, and the other nodes show negligible change. The mean change of state in any node is inversely related to the size of the network; this is because as the size of the network grows, the external network input becomes less significant compared to the influence of other nodes. After plasticity, the same pattern is repeated, with the input node showing a much greater change in state in response to a change in input. However, the overall level of response is much greater than the pre-plasticity networks.

6.5.3.2 Randomly connected networks

Biological neural networks are not fully connected, but are much more sparsely connected. For this reason we studied the effects of homeostatic plasticity on signal propagation in 10-node CTRNNs where connectivity was based on the random graph scheme

⁵There are $N(N + 2)$ parameters in a fully connected N -node CTRNN; in 1-node, 3-node, 5-node and 10-node networks there are 3, 15, 35 and 120 parameters respectively. We realise that the numbers of networks used to generate the data do not therefore reflect an even sampling rate but feel that sufficient data was produced to suffice for the intended demonstration.

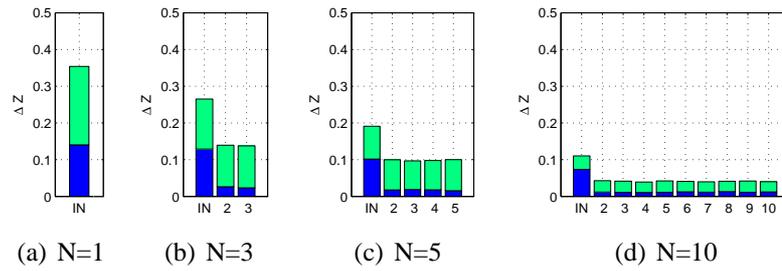


Figure 6.4: Signal propagation in fully connected CTRNNs after homeostatic plasticity has been applied. Mean change in node firing rate in response to a random change in network input is shown for N -node networks for $N \in \{1, 3, 5, 10\}$. Mean changes in node firing rates in response to stimuli are increased by homeostatic plasticity: dark grey represents pre-plasticity level, light grey is post-plasticity increase.

devised by Erdős-Renyi [57, 136]. In these graphs edges between vertices are assigned at random with a fixed probability (see Figure 6.2(b)). Here we created networks by assigning afferent connections between each pair of nodes with fixed probability, generating a random weight value for each connection created. These networks are not intended to mimic the structure of biological neuronal networks (which in any case varies in different species and in different regions of the brain), but simply to give an idea of signal propagation in more sparsely connected networks.

Ensembles of 1000 networks were generated for $P(\text{Connection}) \in \{0.0, 0.1, \dots, 1.0\}$. A single node in each network was designated the input node and received input drawn from the range $[-5.00, 5.00]$, held fixed for 200 timesteps. Signal propagation was measured over 1000 input presentations. Homeostatic plasticity was applied for 500 timesteps and then signal propagation was measured again.

The effect of changes in input on randomly connected networks is shown in Figure 6.5. Input nodes are most affected for all P values, as would be expected, but as connectivity increases the effect of input decreases due to the increased influence of input from other nodes. Hidden nodes are most affected at intermediate connectivity rates of around 20-30%; below this connectivity rate there are insufficient connections for signals to be able to propagate, above this rate the large number of inputs each nodes receives reduces the effective influence of the input node. Homeostatic plasticity significantly increases the effect of input in all cases.

6.5.3.3 Feed-forward networks

An example of the feedforward architecture used is given in Figure 6.2(c). Each node has a self-connection and receives input from every node in the preceding layer; there are no

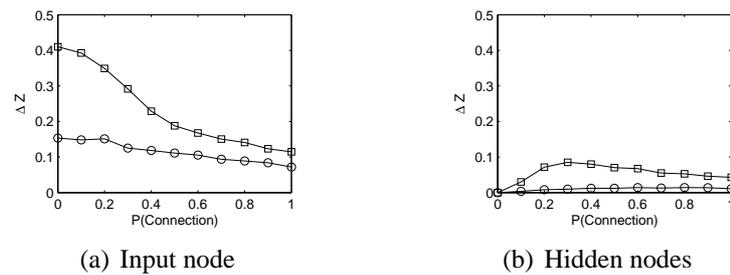


Figure 6.5: Signal propagation in CTRNNs based on Erdős-Renyi random graphs. Networks are created by assigning afferent connections between each pair of nodes with fixed probability. Mean change in node firing rate in response to a random change in network input is shown for 10-node networks for $P(\text{Connection}) \in [0.0, 1.0]$. Mean changes in node firing rates in response to stimuli are increased by homeostatic plasticity: circle and square markers represent pre-plasticity and post-plasticity levels respectively.

lateral connections within a layer or return connections to the preceding layer. This architecture is included in order to look at signal propagation through multiple network layers. While feedforward architectures are in many cases studied without self-connections to the nodes, self-connections were included in the feedforward topology used here to allow easy comparison with the fully connected and randomly connected topologies.

For this architecture input was applied to all nodes in the first layer of the network, and signal propagation was measured as the magnitude of the change in the firing rate vector at each layer subsequent to a change in input. Each node in the input layer received input from the range $[-1.00, 1.00]$ which was held constant for 100 timesteps. The mean size of change was measured over 100 input presentations, for ensembles of 200, 1000 and 2000 networks with 1, 3 and 5 nodes per layer respectively.⁶

Figure 6.6 shows the mean change in state vector at each layer for the ensembles generated, before and after the application of homeostatic plasticity. It can clearly be seen that prior to plasticity, the change in input typically does not affect many downstream layers of the network for any of the network sizes. For the $N = 1$ networks (effectively chains of individual neurons) only neurons in the first 3 layers are affected, and by small amounts. As N rises the signal travels further, but even for $N = 5$ the signal does not get further than the 10th layer. The signal travels further when N is larger because not only is the change in the input vector more significant (incorporating a change in N component dimensions), but each neuron receives input from more neurons in the previous layer (recall that in the feedforward architecture each neuron receives input from all neurons in

⁶The number of parameters in a feedforward network with L layers of N nodes is $LN(N + 3)$; for networks with 25 layers of 1, 3 and 5 nodes there are 100, 450 and 1000 parameters respectively. Obviously the sampling rate is low, but again we feel that sufficient data has been generated to support the argument here.

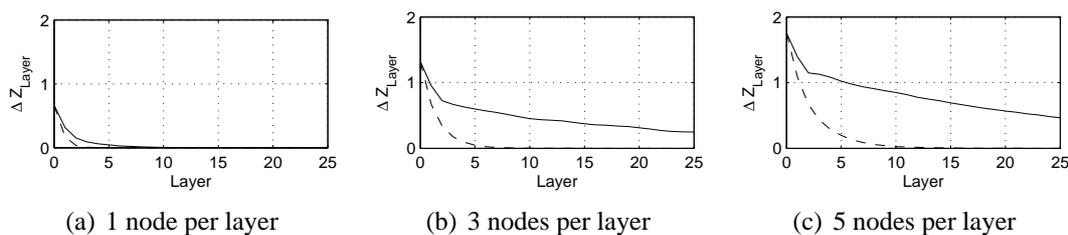


Figure 6.6: Signal propagation in feedforward networks with 25 layers of 1, 3 and 5 nodes, shown as mean change in firing rate vector for each layer caused by a random change in network input. Layer 0 in the plots represents the input vector. Plots shown for networks before (dashed line) and after (solid line) homeostatic plasticity has been applied. Note that ΔZ_{Layer} refers to vector change in firing rate of all nodes in a layer rather than change in firing rate of an individual node.

the previous layer as well as its own self-stimulation; there are no lateral or backwards connections). This has the effect of amplifying the change in state at each layer and thus allows the signal to travel further. It is worth pointing out that if any of these feedforward controllers were used for robot control, before plasticity they would produce a robot that never did anything. Changes in input never cause a change in output, meaning that an agent controlled by the network would never change its behaviour in response to input.

After the homeostatic plasticity has been applied, the change in input clearly causes a change in state at layers much further downstream from the input layer than beforehand. This trend is seen in all the network sizes tested, although it is more prominent in the networks with larger N since these are inherently more conducive to signal propagation because of their greater level of connectivity.

These results suggest that if a feedforward architecture is used for any purpose, homeostatic plasticity should be considered as a method of ensuring that signals are passed usefully from input to output.

6.5.4 Greater likelihood of oscillations

Ensembles of fully connected N -node CTRNNs were generated for $N \in \{3, 5, 8, 10\}$. 500 networks of each size were generated with connection weights randomly drawn from the range $[-10.00, 10.00]$, biases drawn from the range $[-10.00, 10.00]$ and decay constants drawn from $[1.00, 5.00]$. Each network was initialised from 100 randomly sampled positions in its state space. After initialisation the network was updated for 100 timesteps to allow it to settle into a stable equilibrium, after which it was observed for another 100 timesteps to determine whether this equilibrium was a fixed point or an oscillatory

dynamic,⁷ with oscillations detected as described in Section 6.3.2.3.

An oscillation in the network was said to exist if at least one of the nodes showed an oscillation in firing rate.⁸ The proportion of initial conditions that led to oscillatory dynamics was measured for each network and used to calculate a mean value for the whole ensemble of networks of that size. Each network then had homeostatic plasticity applied for a period of 400 timesteps with $\tau_w = 40$ and $\tau_b = 20$, after which the proportion of initial conditions that led to an oscillatory stable state was calculated as before.

Figure 6.7 shows the proportion of initial conditions leading to stable autonomous oscillations in fully connected CTRNNs before and after homeostatic plasticity has been applied. There is a positive correlation between the size of the network and the likelihood of an oscillatory dynamic, reflected in the increased proportion of initial conditions leading to oscillation as network size increases.

Before plasticity is applied, it is noticeable that even in the 10-node networks there is only a small ($\approx 6\%$) proportion of initial conditions leading to autonomous oscillations. This suggests that randomly parameterised CTRNNs are not very likely to display rhythmic dynamics. However, after plasticity has been applied, networks of all sizes are much more likely to oscillate. This can be intuitively interpreted as homeostatic plasticity making networks more sensitive and thus more likely to oscillate. The removal of saturation effects makes sure that all nodes play a part in network dynamics, increasing the number of effective possible oscillatory subcircuits and making the network as a whole more conducive to maintaining autonomous rhythms.

This finding resonates with Mathayomchan and Beer's demonstration that centre-crossing CTRNNs are more likely than standard-form CTRNNs [127] to show oscillations. Oscillations are thought to be important for a variety of motor-control tasks and the results shown here suggest that homeostatic plasticity might be a good method for increasing the likelihood of their occurrence in CTRNNs.

6.6 Agent-level effects of homeostatic plasticity: anecdotal evidence

It is hard to quantify agent behaviour without having a particular task in mind. However, here we will present some evidence that homeostatic plasticity can make the behaviours exhibited by an agent more 'interesting'. Without making any claims as to the usefulness of the properties described here, we wish to draw attention to the increased level of inter-

⁷No distinction was made between limit cycles and chaotic dynamics.

⁸In fact any oscillatory network dynamics must necessarily involve at least 2 nodes.

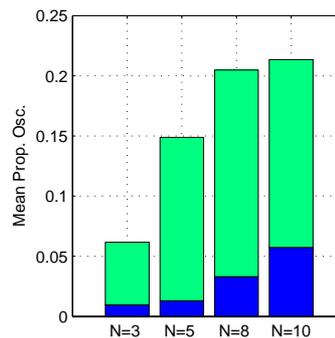


Figure 6.7: Mean proportion of initial conditions that lead to oscillatory autonomous dynamics for fully connected N -node CTRNNs after homeostatic plasticity has been applied (shown for $N \in \{3, 5, 8, 10\}$). There is a significant increase after the plasticity has been applied: blue (dark grey) shows the pre-plasticity level, green (light grey) shows the post-plasticity level.

action between the agent and its environment that occurs when homeostatic plasticity is used.

The robot used is that described in Section 6.3.1 above. We present two sets of trajectory plots describing the motion of the robot was placed in an environment containing light sources. The first set are from runs where the robot was controlled by a randomly parameterised homeostatic plastic CTRNN, and the second set are from runs where the robot was controlled by randomly parameterised non-plastic CTRNNs before and after a period of plasticity.

6.6.1 Robots controlled by homeostatic plastic CTRNNs

Parameter sets for 6-node CTRNN networks and associated sensor/motor gain parameters were randomly generated and instantiated as robot controllers. Figures 6.8(a), 6.8(b), 6.8(c) and 6.8(d) show motion for 4000 time steps (20000 Euler integration steps) where the light source was randomly repositioned every 800 time steps. The symbol \times marks the position of light sources and the symbol $+$ marks the start position of the robot. Note the different scales of the plots.

Robots controlled by randomly generated non-plastic CTRNNs almost always rotated on the spot and ignored the light source. Sometimes they moved in circles, but their motion was consistently a steady rotation. This is explained by the observations of the behaviours of fixed random CTRNNs; they are not generally sensitive to external input because of saturation and tend to give constant output irrespective of input. If the motor neurons fire at a constant rate the differential drive kinematics will produce a constant

angular velocity for the robot, which gives the types of rotational movement observed. An example of this is shown in Figure 6.8(a). The rotation cannot be seen in the plot as the robot was spinning about its axis almost without any movement relative to the light source. The net distance travelled by the robot is very small, and is due to gradual drift as the robot's rotation is very slightly influenced by the different light intensities its sensors experience as it rotates.

In contrast to the non-plastic CTRNNs, robots controlled by the CTRNNs with synaptic scaling displayed quite complex behaviours. Since the synaptic scaling ensured that they were sensitive to external input, these agents were strongly influenced by the light source. Their motion, as seen in Figure 6.8(b), is complex; the easiest description is to liken it to that of a moth near a flame. The robot approaches the source, then veers away, then approaches, then spends some time moving erratically away from the light source. Periods of phototactic behaviour occur, but then cease, as do periods of photo-aversion. The difference in behaviour for different randomly generated networks is considerable, but the plot shown is representative of the general class of behaviours observed.

Interestingly, many of these agents showed occasional periods of phototaxis, and some (roughly 10% by visual inspection) of the agents controlled by CTRNNs with synaptic scaling displayed stable phototaxis. In these cases the typical pattern was an early period of erratic behaviour, while the synaptic scaling configured the network to be sensitive to external input, followed by consistent phototaxis. An example of this is shown in Figure 6.8(c), where the robot approaches and then circles the light sources.

Robots controlled by CTRNNs with adaptive bias also displayed interesting behaviours, which were generally some form of cycloidal motion (see Figure 6.8(d)). The cycloidal motion is caused by the interaction of the adaptive bias with the random set of fixed weights causing oscillations in individual nodes. The cycloidal motion means that the robot moves relative to the light source, in contrast to the near-stationary rotation of the non-plastic CTRNN controllers. The cycloidal motion displayed was generally the same irrespective of how far the agent was from the light source (Figure 6.8(d)), since the adaptive bias can counteract changes in ambient light intensity.

6.6.2 Robots controlled by non-plastic CTRNNs (before/after plasticity)

We have seen that when homeostatic plasticity is active, it creates a greater level of interaction between the robot and the light source stimulus. To show that this increased level of interaction is a feature of increased sensitivity and reduced saturation, and is not caused by the increase in the complexity of the network dynamics, we now show plots of agent

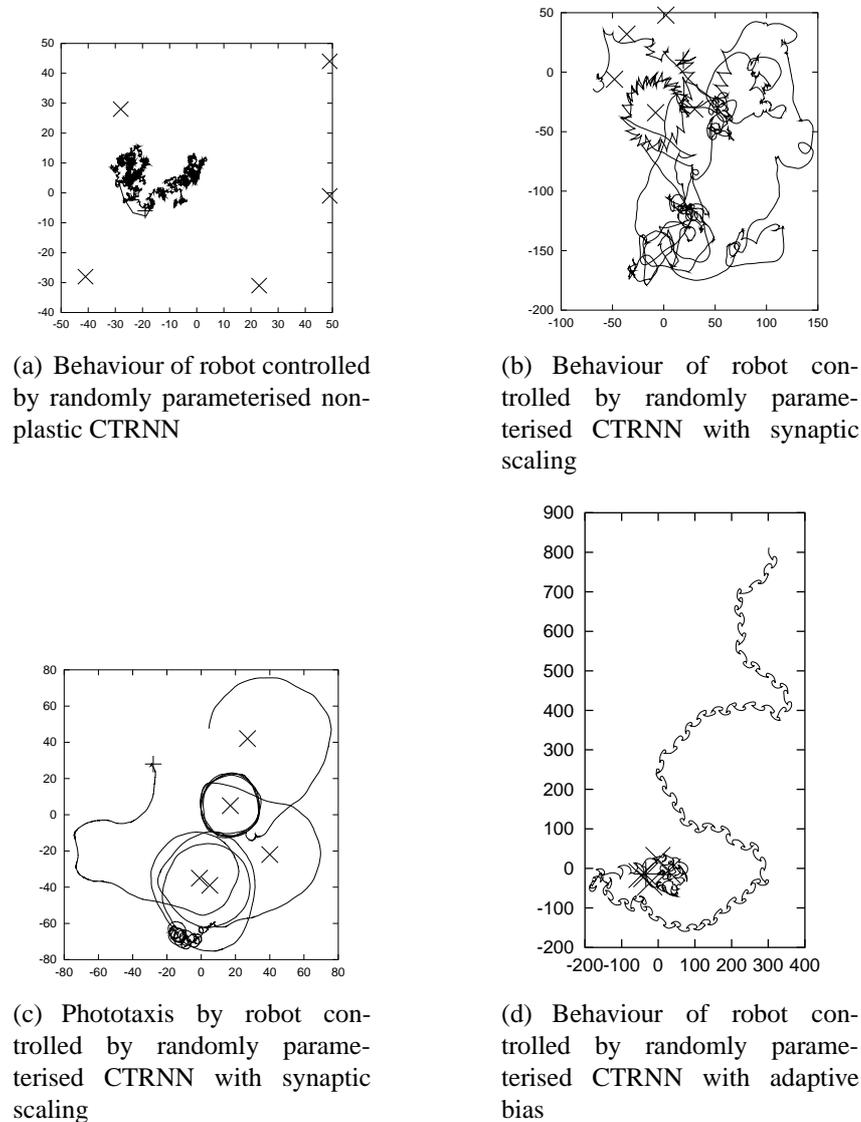


Figure 6.8: Motion plots of photo-sensitive robots controlled by randomly parameterised homeostatic plastic CTRNNs. Robots were placed into an environment containing a single light source which was periodically moved to a new random position relative to the robot. X marks the position of a light source, the line plots the trajectory of the robot's motion over time. Note the different scales of the different sub-plots.

motion before and after a period of plasticity. Here the robot is always controlled by a 6-node non-plastic CTRNN; motion plots are shown for the agent reacting to a series of light sources before and after a period in which homeostatic plasticity was applied to the network.

Figure 6.9 shows the behaviour of the robot for three different randomly parameterised controllers before and after the application of homeostatic plasticity. The motion traces show that initially the agents do not react to the light source. This is because their control networks are largely saturated and they do not change their activity in response to the light falling on their sensors; the controllers give constant output which is translated into stereotyped agent behaviour of rotating on the spot. When the homeostatic plasticity is applied it stops the control networks giving a saturated response so that they now react to the light source. The nature of their interaction with the environment is not of any especially useful form, but an interaction exists.

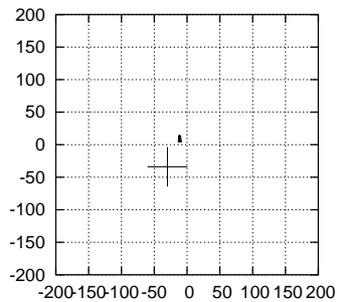
6.7 Summary of results

In this chapter we have looked at the effect of homeostatic plasticity on and in CTRNNs. We have seen that when homeostatic plasticity is active in a CTRNN, it creates the possibility of slow oscillations in a single isolated node due to the interaction of the plastic mechanisms with firing rate. Plasticity also adds a slower timescale to network dynamics, which can mean that networks take much longer to converge to a repeating sequence representing a fixed point or limit cycle attractor. Agents controlled by homeostatic plastic CTRNNs interact much more strongly with their environment.

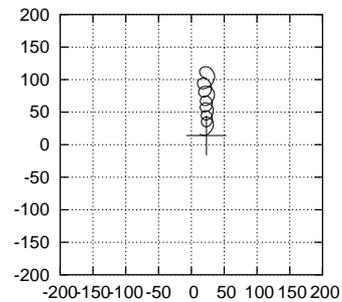
We have also seen that when homeostatic plasticity is applied to a CTRNN and then switched off again, so that the plasticity is used just as a developmental mechanism to condition the network before use, the sensitivity of each node to input increases. This means that signals can propagate much further through a network and that oscillations are much more likely to occur. Agents controlled by CTRNNs that have been conditioned with homeostatic plasticity interact more strongly with their environment.

6.8 Implications for robotics

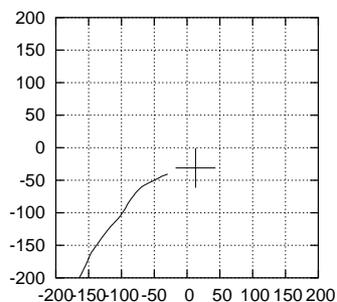
In this chapter we have characterised the effects of homeostatic plasticity in and on CTRNNs. In Chapter 7 we look at the implications of these effects for the evolvability of CTRNNs, but before doing this there are a few observations we can make about the utility of homeostatic plasticity for robotics more generally.



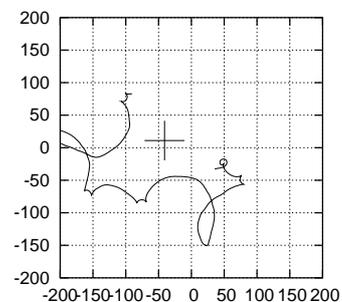
(a) Example 1: Before plasticity



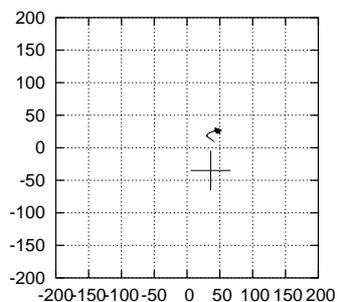
(b) Example 1: After plasticity



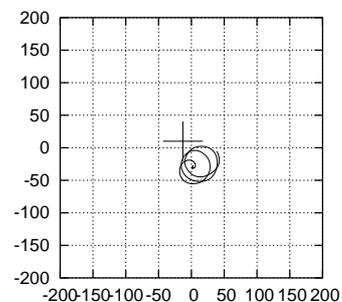
(c) Example 2: Before plasticity



(d) Example 2: After plasticity



(e) Example 3: Before plasticity



(f) Example 3: After plasticity

Figure 6.9: Behaviour examples. Motion of three photosensitive CTRNN-controlled agents before and after homeostatic plasticity is applied. Crosses mark the position of the light sources. Before plasticity is applied the agents do not react to the light source, after plasticity has been applied they are more sensitive to stimuli and there is a stronger level of interaction.

The introduction of an additional slower timescale into network dynamics may allow increased behavioural complexity, and may also have a significant effect on the stability of the system [24]. Homeostatic plasticity performs this function since it operates on a timescale typically an order of magnitude slower than firing rate dynamics.

Node sensitivity is an important basic property in any neural network. If nodes are not sensitive to input then they cannot do any useful information processing, and for this reason homeostatic plasticity may be useful in some circumstances to ensure that an appropriate level of sensitivity is maintained.

Similarly, signal propagation is vital to network information processing. If a change in input does not cause a change in network output then no computation can be performed. In fully connected network architectures the path length from input to output is only one step, but in multi-layer networks signal propagation becomes a more significant consideration. As demonstrated by the results shown above for feedforward architectures, randomly parameterised non-plastic CTRNNs typically do not pass any signal more than a few layers downstream. Homeostatic plasticity increases the distance a signal is propagated by making individual nodes more sensitive. Thus, homeostatic plasticity may be useful for improving signal propagation in multi-layer networks.

Oscillations in network dynamics are thought to be important for a variety of motor control tasks in biology, and by extension in robotics also. Homeostatic plasticity not only makes network-level oscillations more likely by increasing sensitivity, but also allows the possibility of slow oscillations in single nodes, and may therefore be a useful mechanism to include in control networks for certain motor tasks. Homeostatic plasticity also has the potentially useful property of adding a slower timescale of oscillation to network dynamics.

The work presented here has been done largely from the perspective of evolutionary robotics, but the properties described above are for homeostatic plastic networks in general and are valid independent of the method used to parameterise the networks.

6.9 Implications for neuroscience

Homeostatic plasticity in biological nervous systems is thought to regulate neural firing around some set point. It has often been ascribed a stabilising role in neural dynamics, and has been suggested as a counter-mechanism to the destabilising positive feedback effects associated with Hebbian plasticity.

However, the results presented here suggest that the network-level result of neuron-level regulation of firing rates is actually increased levels of activity passing between nodes. Homeostatic plasticity makes each node more responsive to input, meaning that

signals propagate further and with greater amplitude. The result is an increased chance of oscillatory dynamics and greater interaction between the network and its sensorimotor environment.

Thus it seems that far from stabilising patterns of neural activity, homeostatic plasticity may actually cause greater excitability at the network level and thereby destabilise network dynamics. It is unwise to make too bold a claim about biological brains in reference to results derived from a grossly simplified cybernetic abstraction, but the results we have obtained here suggest that the current neuroscientific view of homeostatic plasticity as a stabilising force may not be true at the network level. The idea that homeostatic plasticity counteracts Hebbian plasticity may, however, still be valid.

Hebbian rules in CTRNNs tend to cause positive feedbacks and lead to node saturation as a result of extreme weight values. Homeostatic plasticity offers a mechanism that can prevent node saturation in CTRNNs, maintaining network sensitivity and promoting integrated dynamics. It may play a similar role in biological nervous systems.

Chapter 7

Does homeostatic plasticity make CTRNNs more evolvable?

7.1 Overview

In this chapter we explore the impact of homeostatic plasticity on the evolvability of CTRNNs for robot control. The chapter begins with some discussion of evolvability and the reasons why an investigation into the evolvability of homeostatic plastic CTRNNs is worthwhile. The experimental method is then described, followed by a presentation of the results achieved. The chapter concludes with some discussion of the results and their implications for evolutionary robotics and neuroscience.

7.2 Homeostatic plastic CTRNNs as an evolutionary substrate

In Chapter 6 we saw that homeostatic plasticity makes CTRNNs more sensitive to input, allows signals to propagate further, increases the likelihood of oscillatory dynamics, and strengthens the interaction between an agent and its environment. Homeostatic plastic CTRNNs are *behaviour-rich* and *poised to behave*. We have also noted in Chapter 3 the findings of Mathayomchan and Beer [127] that centre-crossing CTRNNs evolved good pattern-generation circuits more quickly than standard CTRNNs [127]. We went on to

show in Chapter 5 that homeostatic plasticity moves CTRNNs towards the centre-crossing condition. If the inclusion of homeostatic plasticity makes CTRNNs behaviour-rich and at the same time makes them similar to centre-crossing networks, then we might expect that it will make CTRNNs more evolvable. However, there are also reasons why this may not be so. As mentioned previously homeostatic plasticity increases the dimensionality of network dynamics by turning network parameters into variables. This may increase the amount of divergence and noise in the dynamics and increase their complexity, which may make consistent performance more difficult to achieve. Also, we have seen that homeostatic plasticity causes a large increase in the time taken for a network to reach an attractor from initialisation, which might mean that networks perform badly in fitness trials unless the trials are of sufficient duration to nullify possible poor performance during the transient phase.

There are arguments for and against an evolvability payoff from homeostatic plasticity, and it is not clear whether it will be useful. The remainder of this chapter will attempt to shed some light on this issue.

7.3 Measuring evolvability

By common usage in evolutionary robotics, *evolvability* refers to how easy it is to optimise a neural controller using a genetic algorithm. We say that a type of controller is more evolvable if it reaches a higher level of performance than other types, or reaches an equivalent level in fewer generations. This is a reasonable common sense definition and one which we will make use of in this chapter, but there are a few complexities of which we must be aware.

When the evolvability of a system is discussed, it is usually in reference to a particular optimisation task and with a particular flavour of genetic algorithm in mind. Often a type of controller is cited as being highly evolvable after it has performed well in a single evolutionary task. This makes the often unstated assumption that there is some measure of generality to evolvability, i.e., that if a class of controller is more evolvable using genetic algorithm A on task X, it will also be more evolvable when optimised with genetic algorithm B to perform task Y. However, there is little evidence for this and it seems likely that evolvability is highly contingent on the form of the genetic algorithm and the specific task it is applied to, as well as the nature of the controller. Good evolvability is likely to result from a synergy between controller, task and genetic algorithm. This may make testing the *general* evolvability of a type of controller difficult.

However, using a similar argument to that in Section 3.4.3, we can claim that while there may be no universal best-evolving controller, it is likely that there will at least be

some controllers that are more evolvable than others with more varieties of genetic algorithm and for more classes of behavioural task. Thus we can justifiably seek to find more evolvable controllers that will improve the methodology of evolutionary robotics and lead to more rapid and robust design.

How should we find these evolvable controllers? To avoid the problem of task specificity highlighted above and try to gain the most general measure of the evolvability of a class of controller, one method is to perform evolvability tests of the controller with many different genetic algorithms and many different tasks. In the absence of a clear theory of evolvability, it may be that this appeal to weight of data is in fact the best approach. However, this method rapidly becomes intractable. Evolutionary robotics experiments are computationally expensive and time consuming, especially if performed on hardware rather than simulation. There is a clear trade-off between generality and tractability.

In an ideal world it would be beneficial to run a large variety of experiments to clearly elucidate the evolvability benefits (or not) of homeostatic plasticity in CTRNNs. However, constraints of time and space mean that this is not possible in this thesis and we will have to be content with the more limited set of experiments offered below. We accept that this is not a complete account, but hope that it will shed at least some light on the issue and perhaps point the way for future work.

The results presented in this chapter should be seen as the first few data points found in a larger study of the evolvability of neural controllers with homeostatic plasticity. It is hoped that other researchers may be tempted to run different experiments with different implementations of the basic theory. Doing so will test the assumptions made in the current implementation and thus allow for greater generality in the construction of a coherent theory.

Here we will make use of the common sense definition of evolvability described above and look for comparative fitness performance achieved with different plasticity schemes and tasks when CTRNN controllers are evolved for simple agent-based behavioural tasks.

7.4 Method

We will look at evolvability of CTRNNs with and without homeostatic plasticity on two simple agent-based tasks derived from original work by Beer [10]. These tasks involve a simulated agent that uses visual information to respond to shapes that fall towards it in a vertical plane. By moving horizontally towards or away from a falling object, the agent can either catch shapes or avoid them. We define two behavioural tasks. The first task is to use visual information to direct movement so that the agent catches all falling shapes. The second task is to discriminate between different types of shape by catching some and

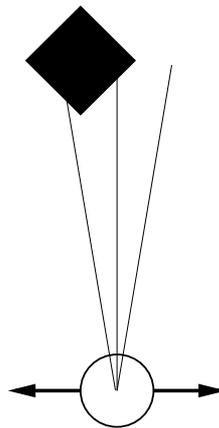


Figure 7.1: The agent has a circular body and is equipped with 3 ray sensors and 2 motors that allow it to move horizontally in response to sensory input.

avoiding others.

Neural controllers are evolved using a standard genetic algorithm. Homeostatic plasticity is applied in various schemes to look at the evolutionary performance of networks incorporating adaptive bias, synaptic scaling, or both mechanisms together. The effect of a pre-trial developmental phase of plasticity is also examined. In all experiments, performance is compared against that of standard non-plastic CTRNNs, which act as an experimental control.

7.4.1 Agent

The agent-shape system (see Figure 7.1) is an idealised abstraction of a situation in which the shapes fall towards a robot that moves horizontally along a straight line on the ground. The agent and the falling shapes are constrained to move in the same vertical plane so that the world is two-dimensional. The agent has a circular body of radius 5 units.

7.4.1.1 Control network

The agent is controlled by a fully connected 5-node CTRNN with 3 sensor neurons and 2 motor neurons. There are no interneurons. The CTRNNs and homeostatic plastic mechanisms are implemented as described above in Sections 3.4.1 and 4.4 respectively, with the homeostatic target range set to be $z \in [0.2, 0.8]$. The network is updated using Euler's forward method with step-size of 0.2 timesteps.

7.4.1.2 Sensors

There are 3 ray sensors mounted on the periphery of the agent and arranged in an upward-facing fan (see Figure 7.1), with each sensor giving a signal presented to the network as the external input (I in Equation 3.1) to a unique node in the CTRNN controller. Each sensor returns a signal proportional to the proximity of any surface with which its ray intersects as specified by Equation 7.1. Sensors may be thought of as similar to laser range-finders. The sensor array spans an arc of $\frac{\pi}{6}$ radians.

$$S = S_{max} \left(\frac{D_{max} - D}{D_{max}} \right) \quad (7.1)$$

where S is the signal returned by the sensor, S_{max} is the maximum sensor value, D is the distance from the ray's origin to the intersected surface, and D_{max} is the maximum distance at which sensors can detect an object. In the experiments presented here $S_{max} = 5$ and $D_{max} = 100$. Sensors only return intensity values related to the distance to the intersection, and encode no other information.

7.4.1.3 Motors

The agent can move left and right horizontally as if mounted on a frictionless rail, with motion determined by the output of the remaining non-sensor nodes. The output of the motor nodes is mapped to the output of two motors (acting in opposite directions) that give an immediate horizontal velocity proportional to the difference in their activity. The agent is assumed to have zero mass, so that there is no momentum or inertia. Agent movement is calculated using Equation 7.2.

$$\tau_x \dot{x} = z_{right} - z_{left} \quad (7.2)$$

where x is the position of the agent, z_{left} and z_{right} are the firing rates of the motor neurons for the left (negative x) and right (positive x) directions respectively, and $\tau_x = 0.2$ is a time constant for the velocity. Since τ_x is for simplicity chosen to be the same as the integration step for the control network, the change in position at every network update is given by the right-hand side of Equation 7.2.

7.4.2 Falling shapes

There are two types of shape, circles and diamonds, both of radius 10 units (i.e., circles are of radius 10 units, while diamonds are sized such that their vertices would lie on the circumference of a circle of radius 10 units). Shapes begin at a horizontal displacement

from the agent that is randomly chosen from 10 possibilities evenly distributed in the range $[x - 25, x + 25]$, where x is the position of the agent. They appear at a vertical distance of 100 units above the agent, then fall until the lowest point of the shape passes below the level of the uppermost part of the agent, at which point they disappear and are replaced by a new shape randomly placed as before.

Since the agent can only sense objects through its 3 ray sensors, the agent's perception of the falling shapes is a continuously varying triple of sensor values that varies as the relative position of the shape from the agent changes. The values returned are also determined by the relief of the leading edge of the shape, so the relationship between the values in the triple provides information that can in principle be used to differentiate between circles and diamonds.

7.4.3 Genetic algorithm

Each evolutionary run is performed with a population of 50 genotypes evolved for 500 generations. The initial population consists of randomly generated genotypes. Each subsequent generation is created from the fittest members of the preceding generation, selected according to the scheme below.

7.4.3.1 Encoding

For all types of controller, the weights, biases and decay constants for all nodes in the network are encoded on the genotype. The genotype is real-valued with allele values on the range $[-1.00, 1.00]$. These values are linearly mapped to appropriate ranges for the associated phenotypic traits when networks are instantiated from genotypes. The genotype for an N -node network is thus an array of size $N^2 + 2N$, i.e., 35 values in the 5-node case. The allele values for weights and biases are interpreted differently for plastic networks and non-plastic networks; for plastic networks the genotype encodes the initial values for the weight and bias *variables*, while in non-plastic networks the genotype encodes the values for the weight and bias *parameters*.

7.4.3.2 Selection

Elitism is used to preserve the best genotypes from each generation unaltered. Each new generation is formed by passing the 5 best genotypes unchanged from the previous generation and making up the rest of the new population from genotypes chosen by roulette wheel selection.

In roulette wheel selection the previous generation is first of all ranked in fitness order. Then, each member of the population is assigned a number of slots in an imaginary roulette wheel. Fitter genotypes have more slots. In a population size of N there are N slots assigned to the fittest genotype, $N - 1$ to the second fittest, and so on down to the least fit genotype, which gets 1 slot. There are thus $N(N + 1)/2$ slots in the complete wheel.

The new generation is created by randomly choosing a slot (spinning the roulette wheel) and using the associated genotype as a parent for the next generation. Since fitter genotypes have more slots, they are on average more strongly represented in the next generation than less fit genotypes. This process is repeated, with replacement, until sufficient parent genotypes have been chosen to form the new population.

7.4.3.3 Mutation / crossover

No crossover was used, so that the population can be thought of as reproducing asexually. The elite genotypes are passed to the next generation without mutation, but all other genotypes are subjected to point mutation with a 3% chance of mutation at each allele. Two forms of point mutation occurred with equal probability. The first type of mutation was carried out by randomly choosing a new value for the allele from the range $[-1.00, 1.00]$. The second type added to the existing allele value an amount randomly chosen from the range $[-0.50, 0.50]$.

7.4.3.4 Fitness function 1: Ball-catching task

In the ball-catching task a series of circle shapes is dropped and the agent has to move so that it intercepts them. The agent is awarded fitness for moving so that it ‘catches’ all the falling objects, a task which involves locating the object and moving to keep the object in the centre of its field of view until it reaches ground level.

In each trial, 20 circles are dropped from 10 possible start positions with randomly chosen velocity. The horizontal component of shape velocity is drawn from a uniform distribution on the range $[-0.3, 0.3]$, while the vertical component is drawn from the range $[-0.5, -0.2]$.

Fitness is calculated as a combination of scores on two criteria. The first criterion is percentage reduction in the horizontal distance between the centre points of the agent and the falling object. The second criterion is minimising the absolute distance between the agent and the shape at the end of the trial, with the score normalised by comparing the actual distance to the greatest possible distance based on the maximum speed of the agent.

Mean values for both criteria are taken over the 20 object presentations in each trial,

summed with equal weighting and normalised to the range $[0.0, 1.0]$. The trial fitness score in the ball-catching task is therefore given by Equation 7.3.

$$Fitness = \frac{1}{P} \sum_{i=1}^P \frac{1}{2} \left[\phi \left(1 - \frac{S_{final}}{S_{init}} \right) + \left(1 - \frac{S_{final}}{S_{max}} \right) \right]_i \quad (7.3)$$

Here S is the horizontal displacement of the object from the agent. $\phi(\dots)$ is a function that returns zero if the argument is outside the range $[0, 1]$, and otherwise leaves the argument value unchanged, i.e., in this case the function returns zero if $S_{final} > S_{init}$ but otherwise returns the argument value. S_{max} is the maximum achievable horizontal distance between agent and shape is calculated by working out the time taken for each shape to reach the ground (based on its vertical velocity) and multiplying by the maximum horizontal speed of the agent relative to the shape. P is the number of presentations ($P = 20$ used here). The indexed square brackets $[\dots]_i$ refer to the data gathered from the i^{th} object presentation. The fitness for a genotype is taken as the mean value over 10 trials in order to promote consistency in performance.

7.4.3.5 Fitness function 2: Discrimination task

The discrimination task is similar to the ball-catching task except that a variety of circle and diamond shapes are presented, with the task for the agent being to avoid circles and catch diamonds.

In this task the shapes fall vertically at a speed of 1 unit per timestep and have zero horizontal velocity. In each trial, 20 shapes (10 circles, 10 diamonds) are dropped from 10 possible positions.

A shape is deemed to be caught if it reaches the ground at a horizontal distance from the agent less than the combined radii of agent and shape, i.e., if the agent and shape overlap at ground level. Objects that are not caught are said to be avoided. Catching or avoidance is viewed as a behavioural expression of the agent's successful discrimination between objects.

Fitness is the proportion of correct classifications performed. Trial fitness for the discrimination task is given by Equation 7.4.

$$Fitness = \frac{1}{P} \sum_{i=1}^P \gamma(i) \quad (7.4)$$

where $\gamma(i)$ returns 1 or 0 when the correct or incorrect classification respectively is made for presentation i , and $P = 20$ is the total number of object presentations in the trial.

Again the fitness for a genotype is calculated as the mean value over 10 trials in order to promote consistency in performance.

7.5 Results

We will present results here from four experiments, in order to tell a comprehensible story about the effect of homeostatic plasticity on the evolvability of CTRNN robot controllers. First we will perform a basic comparison of evolvability of CTRNNs with different plasticity schemes. Next we will look at the effect of adding a developmental period to the evolutionary trials, then at how homeostatic plastic networks perform next to centre-crossing networks. Finally we will consider the performance of larger non-plastic CTRNNs, in order to determine the effect on evolvability of greater numbers of degrees of freedom in network dynamics.

7.5.1 Experiment 1: Comparative evolvability of CTRNNs with different homeostatic plasticity schemes

For this set of experiments CTRNN controllers with different plasticity schemes were evolved using the method and tasks described above. Ten evolutionary runs were performed for each scheme to make sure that the results gained were representative; while 10 measurements is not a large sample, it should be sufficient to reduce the influence of occasional anomalous data whilst remaining within the bounds of tractability.

The plasticity schemes used are given below. All networks had 5 nodes and were fully connected, with sensor/motor connections as described in the method above.

- Non-plastic CTRNN
- CTRNN with synaptic scaling
- CTRNN with adaptive bias
- CTRNN with both synaptic scaling and adaptive bias

Each type of network was tested on both the ball-catching and discrimination tasks. Figure 7.2 shows the average evolutionary performance of controllers with each plasticity scheme on the ball-catching task, taken over 10 evolutionary runs. Figure 7.3 shows the performance in each run individually for comparison. Figures 7.4 and 7.5 show the same results for the discrimination task.

The results for the ball-catching task show that the best individual performances were by non-plastic networks, but that the population average fitness was generally slightly higher for the plastic controllers (Figure 7.2). All of the plastic controllers reach a high level of fitness in fewer generations than the non-plastic controllers, with the average fitness rising more quickly in the early stages of the evolutionary runs.

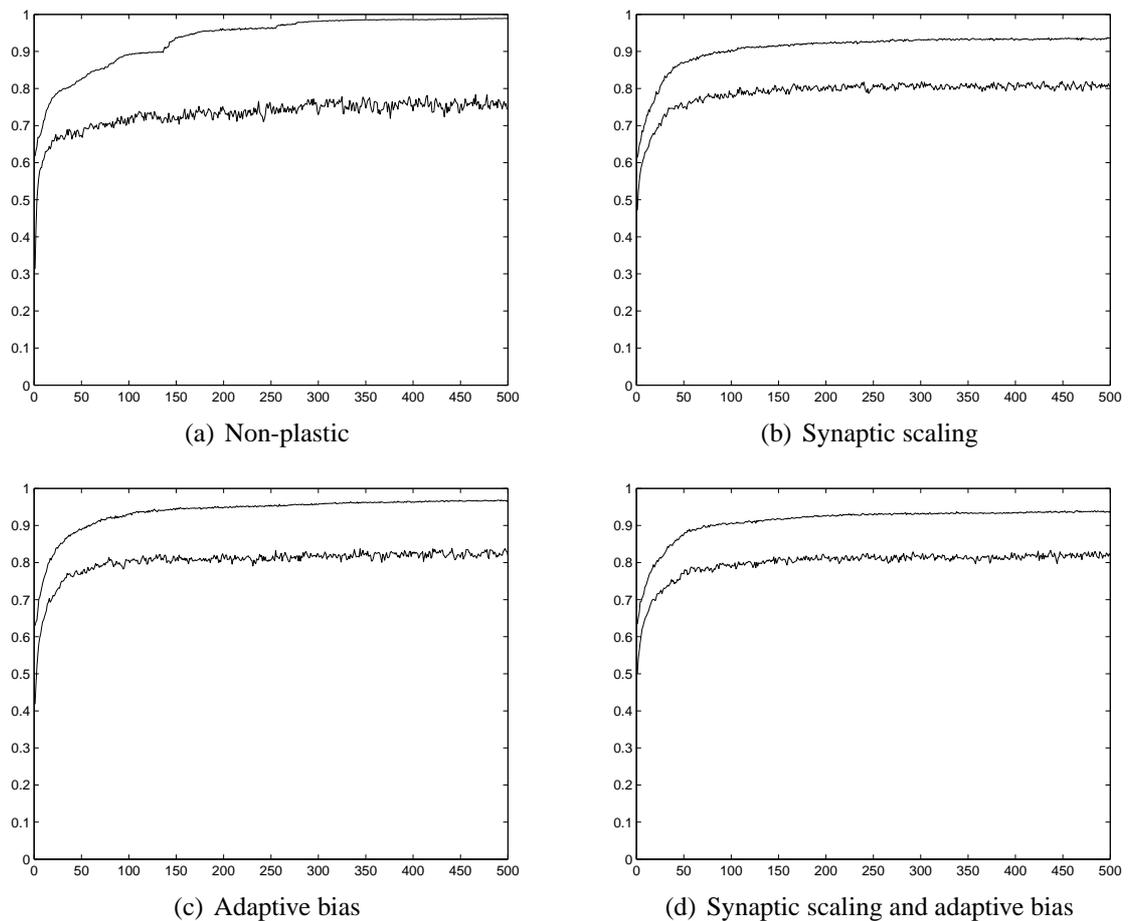


Figure 7.2: *Ball-catching task*. Plots show mean performance over 10 evolutionary runs of CTRNNs with different homeostatic plasticity schemes. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. Non-plastic networks reach the highest levels of fitness by an individual in the population, but plastic networks are more consistent, with a higher population mean fitness.

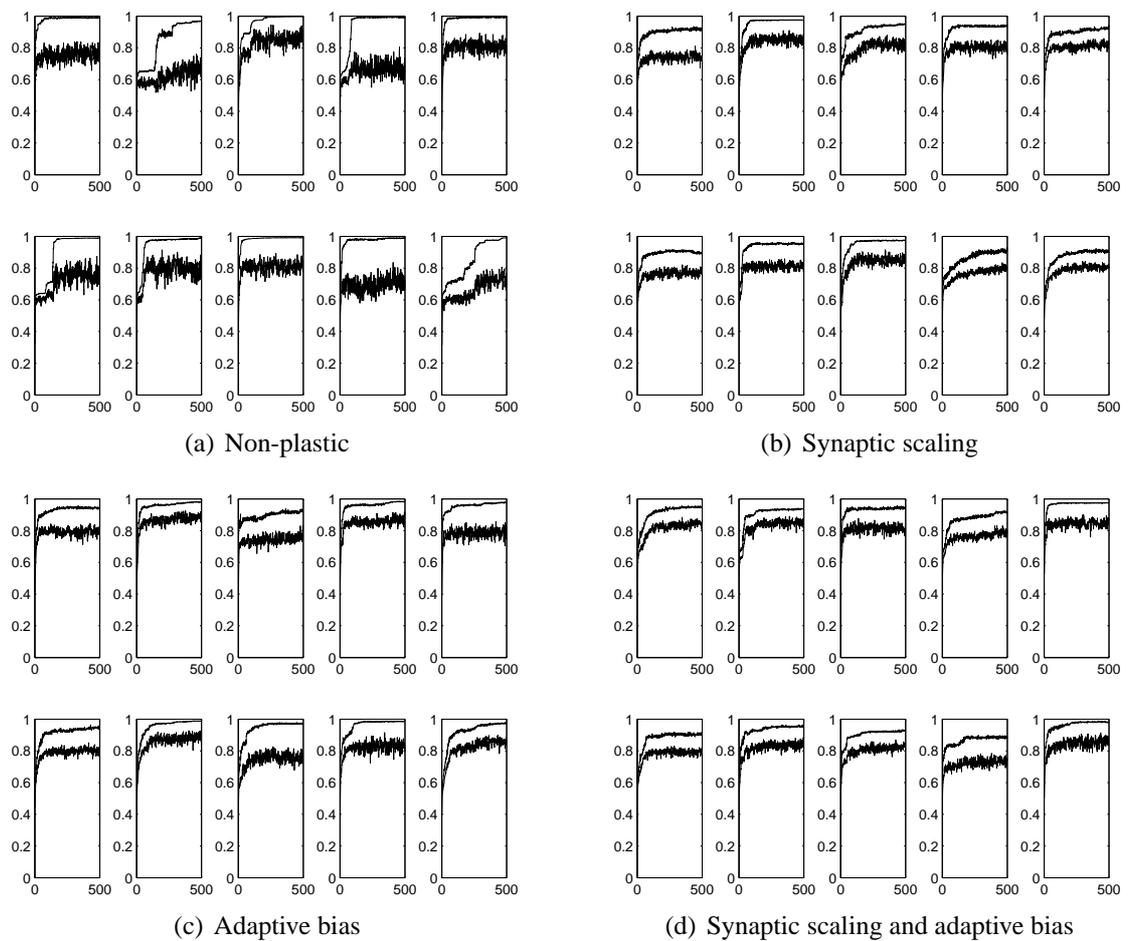


Figure 7.3: *Ball-catching task*. Plots show individual run performances (for 10 evolutionary runs) of CTRNNs with different homeostatic plasticity schemes. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. Non-plastic controllers show jerky increases in fitness during the early stages of the evolutionary runs, while plastic controllers show a smoother increase.

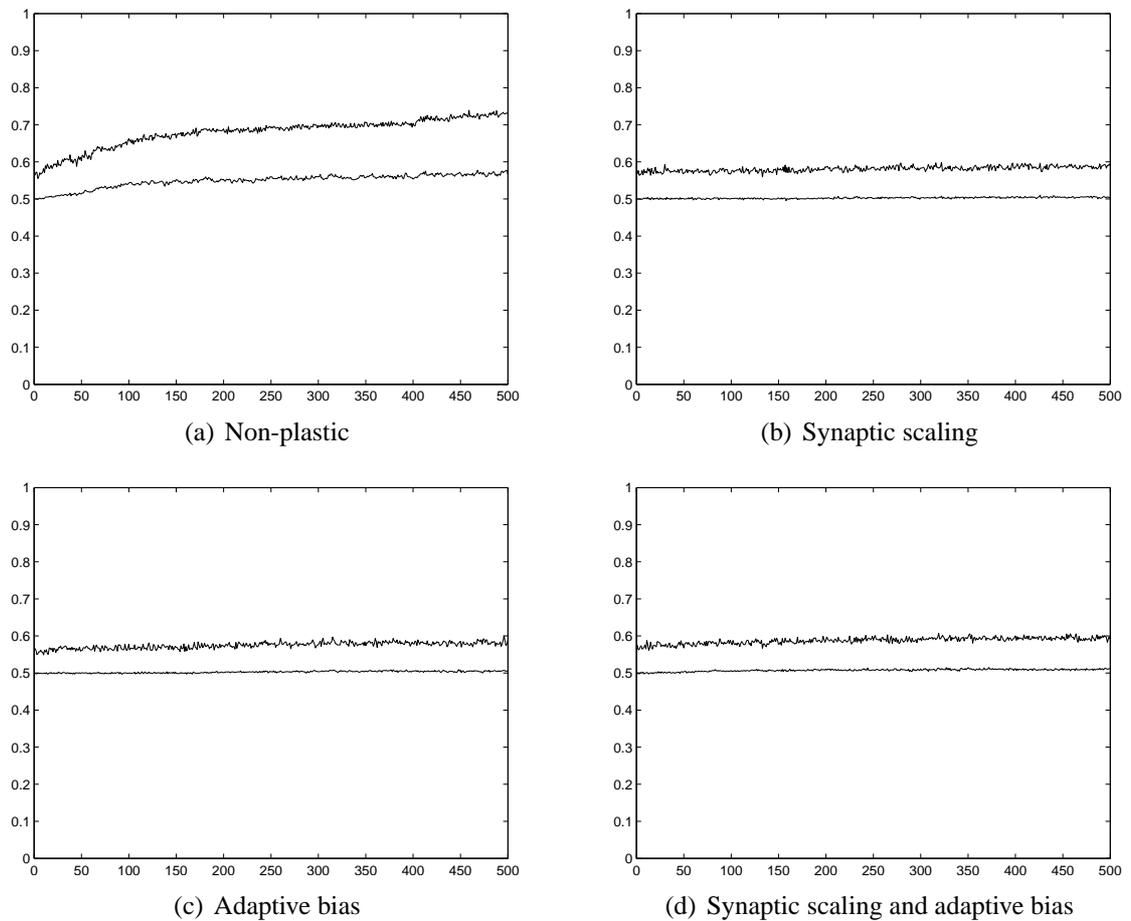


Figure 7.4: *Discrimination task*. Plots show mean performance over 10 evolutionary runs of CTRNNs with different homeostatic plasticity schemes. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. All kinds of network perform poorly, but non-plastic networks are slightly better than the plastic networks.

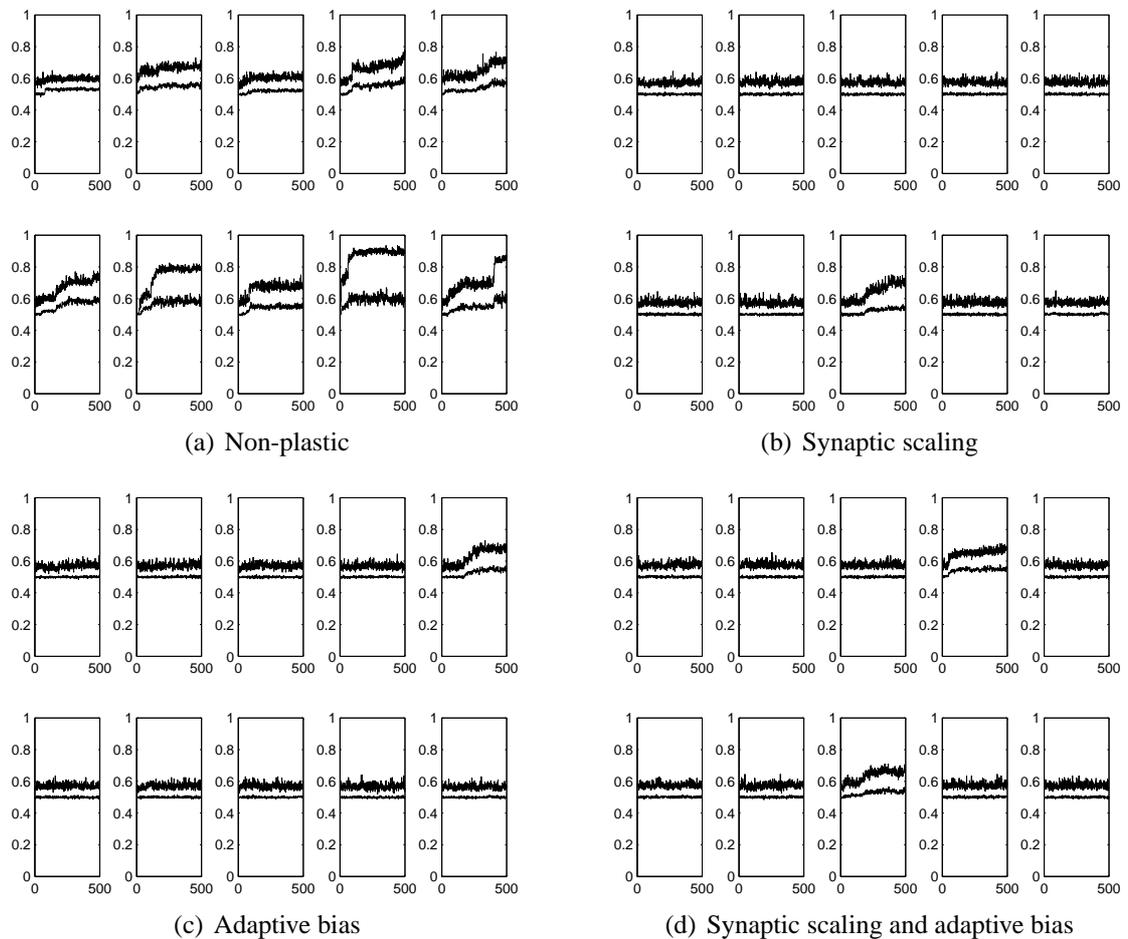


Figure 7.5: *Discrimination task*. Plots show individual run performances (for 10 evolutionary runs) of CTRNNs with different homeostatic plasticity schemes. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. Most of evolutionary runs for the non-plastic networks show some degree of progress, but only a few of the plastic networks show any improvement in evolutionary time.

The plastic controllers also all show a smooth increase in fitness, whereas the non-plastic networks exhibit a more jerky increase. While this jerky increase would be smoothed out if the sample size of runs was increased, it is notable that the plastic controllers are a lot more consistent than the non-plastic networks, showing greater similarity in their performances in different runs than the non-plastic networks (see Figure 7.3). The plastic controllers always quickly evolve a reasonable level of fitness, while the non-plastic controllers improve slowly to begin with but go on to eventually reach a higher best performance level. There is little difference between the different plastic networks, with all giving a similar level of performance.

The results for the discrimination task show that it is a more difficult evolutionary challenge than the ball-catching task. Figure 7.4 shows that the non-plastic networks achieve only a moderate level of fitness, while the plastic networks universally perform poorly.

It is possible to get a fitness score of 0.5 in the discrimination task by shooting off in one direction, hence catching nothing and thus making 50% correct classifications. This strategy is likely to be present in the initial population of all of the controller types. The plastic networks appear never to improve on this strategy, with the best networks in each generation probably just being lucky enough to make a few additional correct classifications accidentally. Looking at Figure 7.5 we can see that the results for the individual runs show that in only a few cases do the plastic networks show any evolutionary progress, and also that this may take several hundred generations of genetic drift to appear.

Comparing the results for both tasks, we see that homeostatic plasticity offers advantages of consistency and quick early progress in the ball-catching task, but is eventually out-performed by the best non-plastic networks in most runs. In the discrimination task we observe that the plastic networks are clearly out-performed by the non-plastic networks. It is not immediately clear why this should be so, though looking at the strategies of the best non-plastic networks suggests one possibility. The best evolved non-plastic networks use an active scanning motion where the agent moves rapidly from side to side as the object approaches. Catching is then performed by narrowing the amplitude of this oscillation, and avoidance by expanding it at the last moment. There is only a small observable difference between the catch and avoid behaviours and it may be that the increased dimensionality of the dynamics in the plastic networks makes them too clumsy and unreliable for this level of sensitive control.

The level of performance of the different homeostatic plastic mechanisms seems to be similar in all regimes. The adaptive bias mechanism has a slight advantage in the ball-catching task, but this is too slight for us to conclude it is better than synaptic scaling in general, especially given our earlier remarks about the contingent nature of evolvability.

The plastic networks also show similar poor levels of performance on the discrimination task.

From this experiment we have seen that in terms of achieving the highest level of fitness the plastic networks perform worse than the non-plastic networks. However, on the ball-catching task the plastic networks offer advantages of quicker early progress in the evolutionary run and more consistency in achieving a reasonable level of performance in multiple runs.

7.5.2 Experiment 2: The effect of a developmental period

Here we want to look at the effect on evolvability of a developmental phase, where homeostatic plasticity is applied for a period in each trial prior to fitness assessment. We looked at the same types of plastic controller as in Experiment 1, updating every network for 6000 timesteps before each trial began. In one set of evolutionary runs the plastic mechanisms were left active during the trial (Figures 7.6 and 7.7), while another identical set of runs was performed where the plasticity was switched off after development for the duration of the fitness trial (Figures 7.8 and 7.9). In the latter case the plasticity is used purely as a mechanism for pre-conditioning networks that have no active plasticity during their functional lifetime.

Figures 7.6 and 7.7 show the performance of the plastic networks with and without a developmental period for the ball-catching and discrimination tasks respectively. The plots for performance without development have been repeated from Figures 7.2 and 7.4 for ease of comparison.

Performance is broadly very similar on first inspection, with similar levels of fitness being reached irrespective of the presence/absence of a developmental period. However, careful examination of performance in the early stages of the evolutionary runs (generations 0-150) for the ball-catching task shows that the developmental period has the effect of reducing the number of generations taken until a high level of fitness is achieved. This effect seems to be most pronounced in the plots for the synaptic scaling and adaptive bias mechanisms acting individually.

For the discrimination task, the developmental period causes a small increase in the levels of fitness reached, but overall performance by the plastic networks on this task remains very poor.

The effect of the developmental period when used as a pre-conditioning phase for non-plastic networks is quite large in the ball-catching task. Figures 7.8 and 7.9 show the results for runs where homeostatic plasticity was applied for a period prior to each fitness trial, but then switched off while the trial was performed.

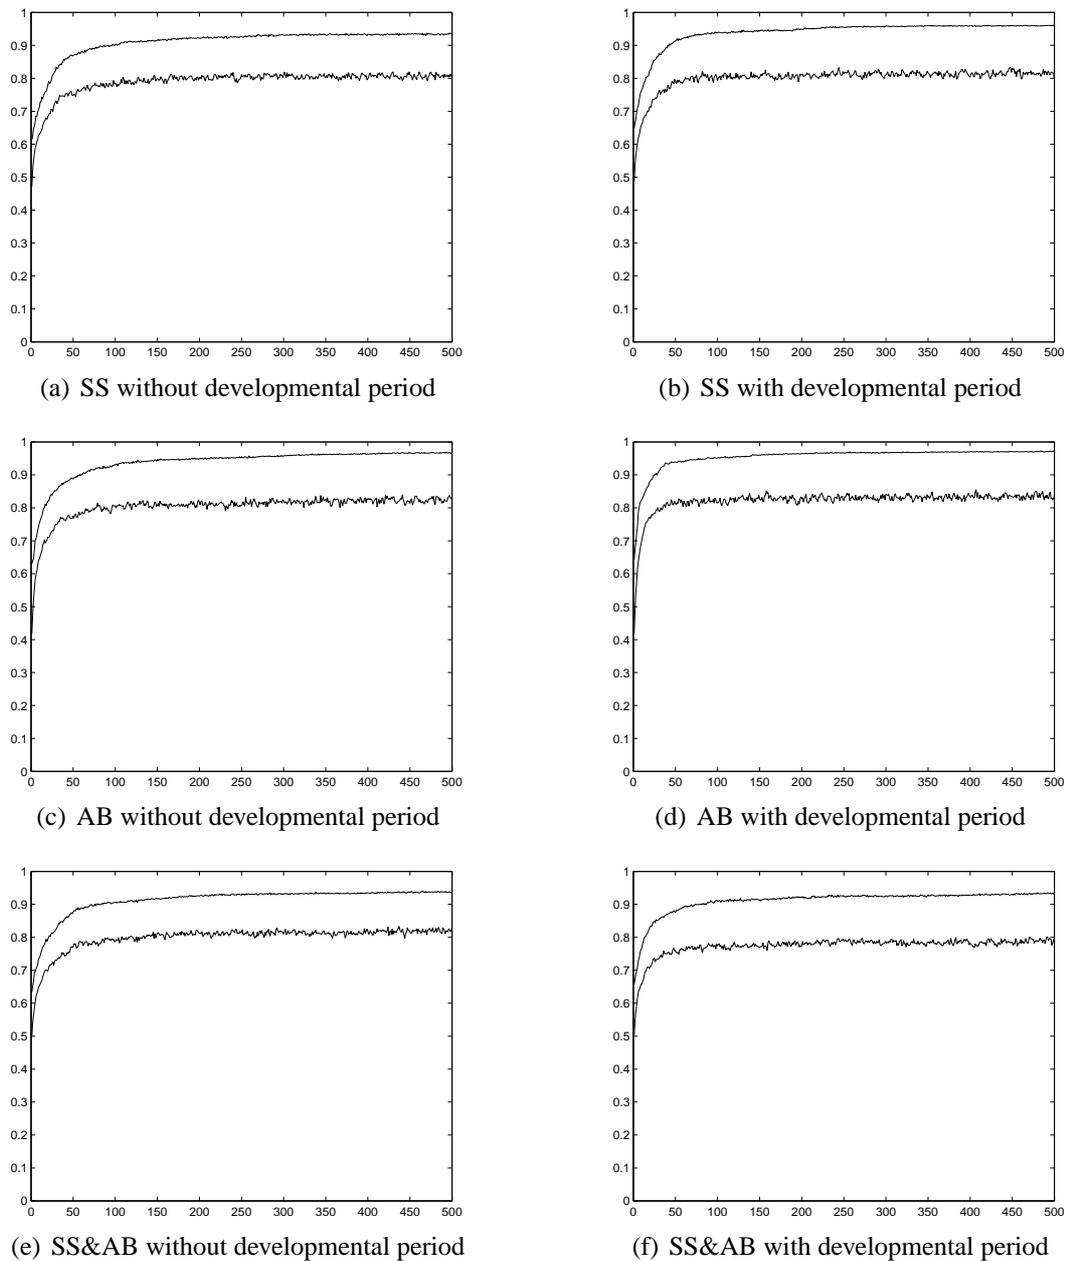


Figure 7.6: *Ball-catching task*. Plots show mean performance over 10 evolutionary runs of CTRNNs with different homeostatic plasticity schemes with and without a pre-trial developmental period. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. SS - synaptic scaling, AB - adaptive bias. The developmental period improves progress in early stages of the evolutionary runs compared to networks without a developmental period, as demonstrated by the quicker attainment of high fitness in the runs with development.

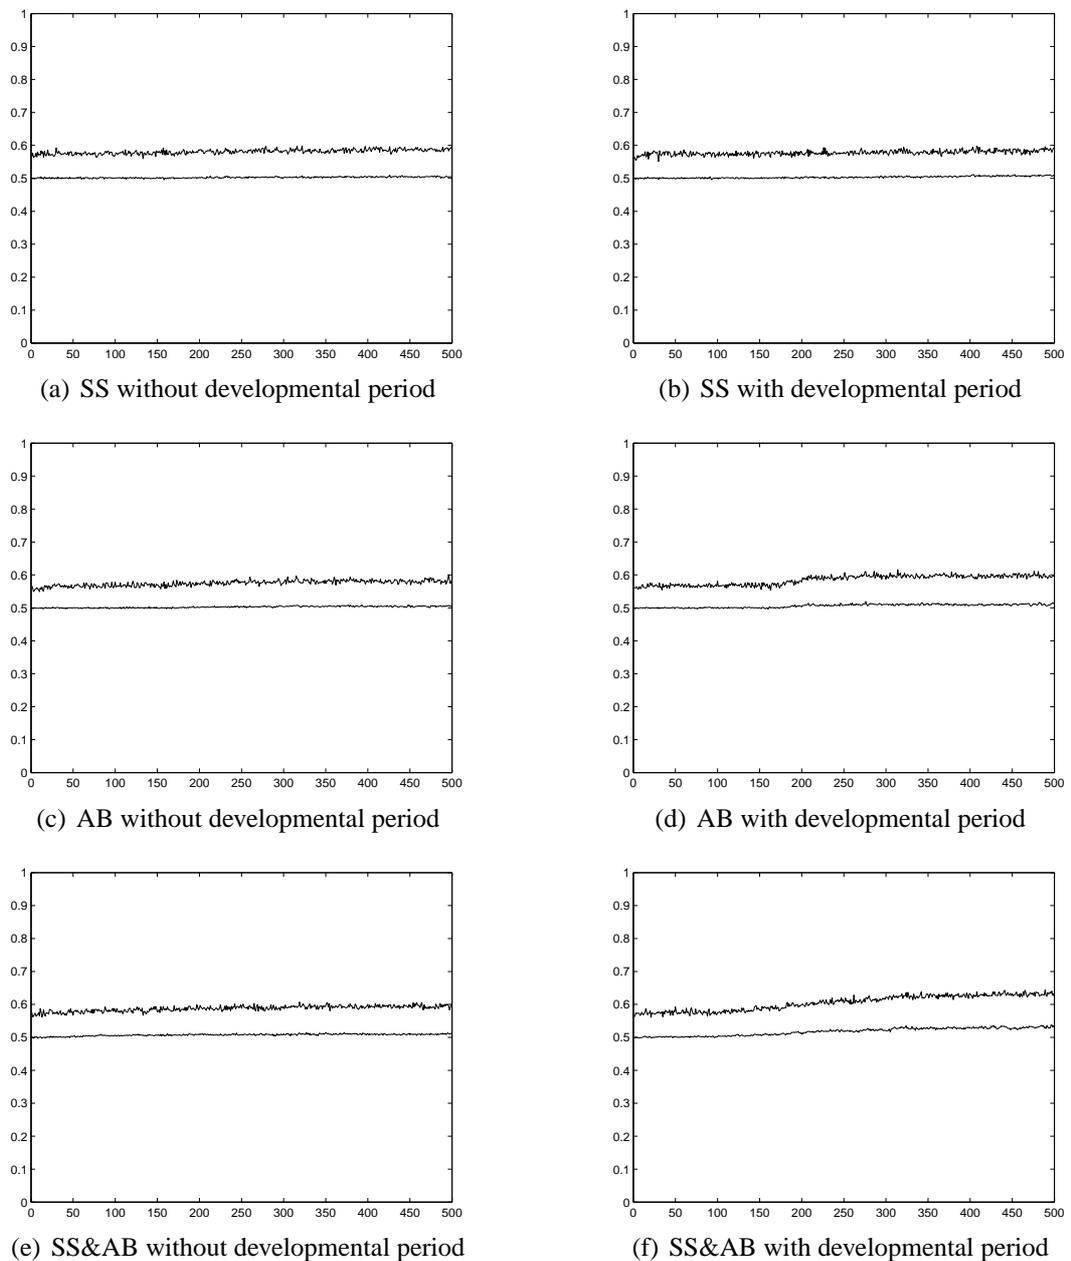


Figure 7.7: *Discrimination task*. Plots show mean performance over 10 evolutionary runs of CTRNNs with different homeostatic plasticity schemes with and without a pre-trial developmental period. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. SS - synaptic scaling, AB - adaptive bias. The use of a developmental period causes a slight improvement in performance but overall levels of fitness are poor both with and without a developmental period.

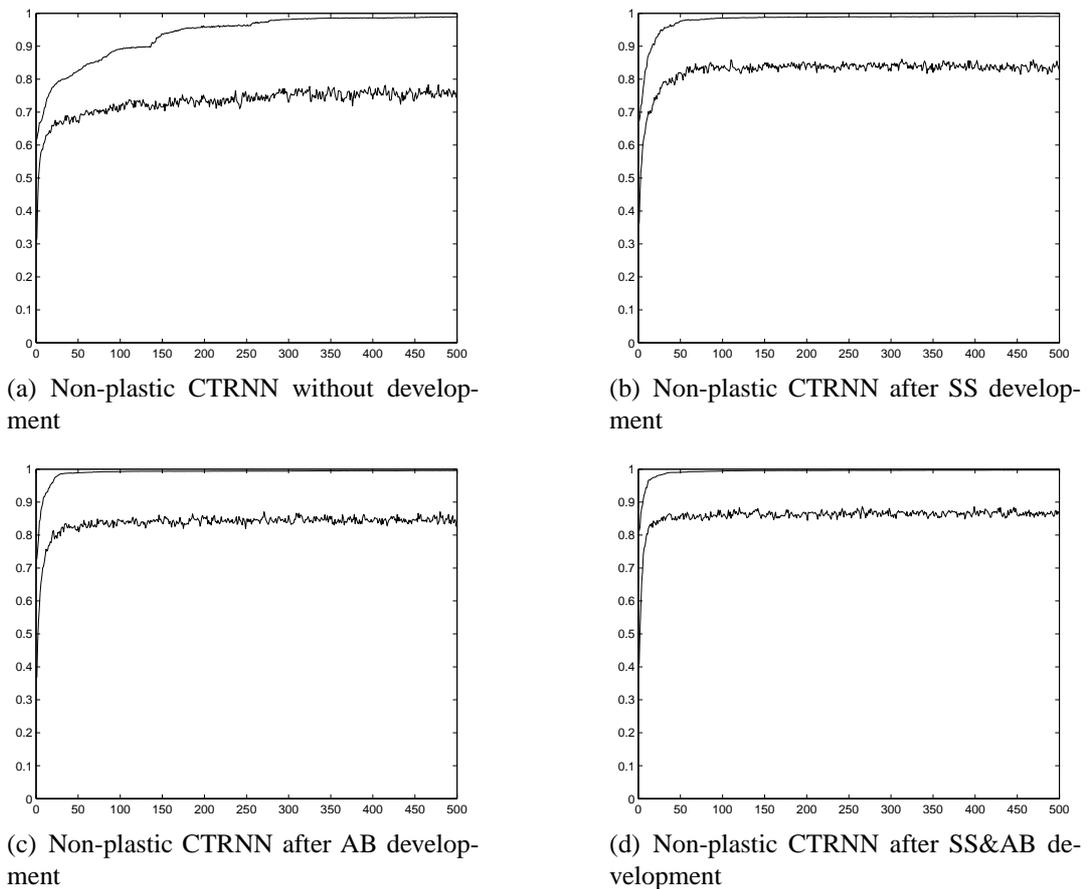


Figure 7.8: *Ball-catching task*. Plots show mean performance over 10 evolutionary runs of non-plastic CTRNNs where fitness trials were performed after a period of homeostatic plastic development. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. Development performed with different plasticity schemes: SS - synaptic scaling, AB - adaptive bias. Using homeostatic plasticity as a developmental mechanism to precondition non-plastic CTRNNs gives a clear advantage in evolutionary performance.

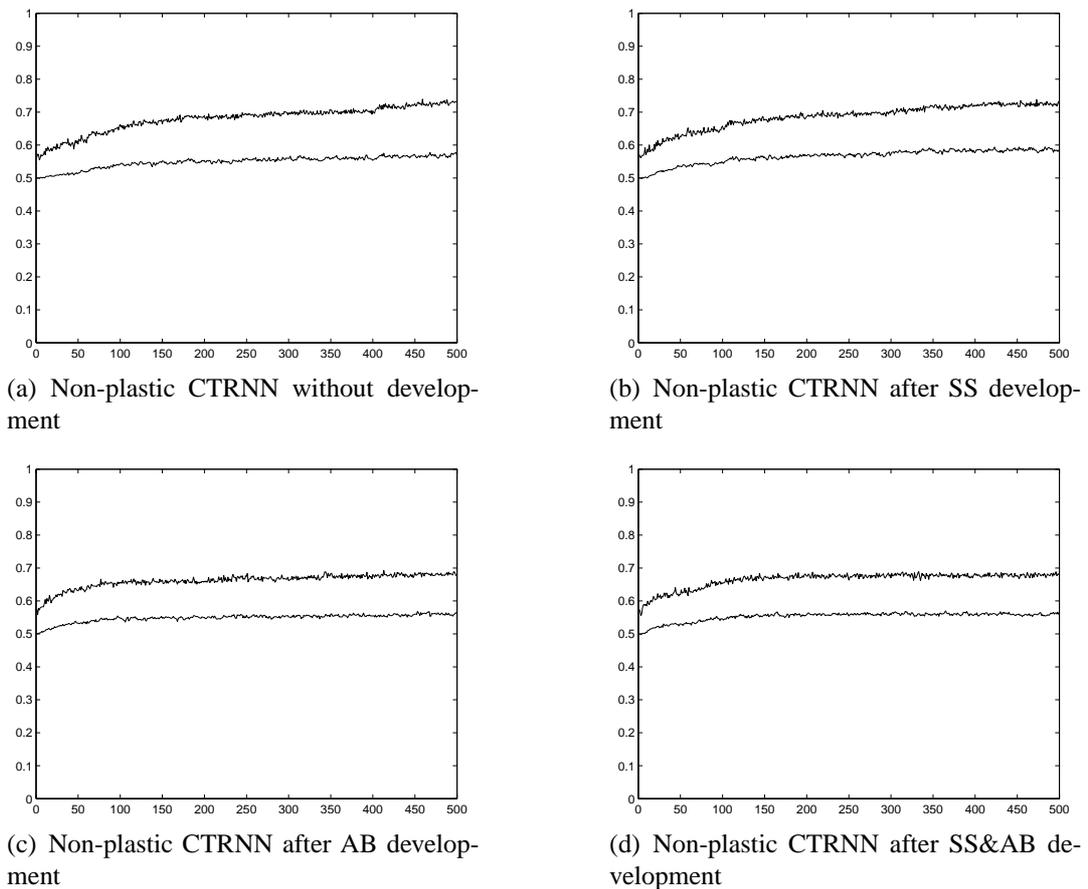


Figure 7.9: *Discrimination task*. Plots show mean performance over 10 evolutionary runs of non-plastic CTRNNs where fitness trials were performed after a period of homeostatic plastic development. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. Development performed with different plasticity schemes: SS - synaptic scaling, AB - adaptive bias. In this task, conditioning non-plastic networks with a period of homeostatic plastic development made no significant difference to performance.

The plots for the ball-catching task (Figure 7.8) show a distinct improvement in the speed of evolution of good controllers, though again the final levels of fitness reached are similar irrespective of development. High levels of fitness on the ball-catching task are reached several hundred generations earlier when non-plastic networks are pre-conditioned with homeostatic plastic development.

Developmental pre-conditioning has very little effect on performance in the discrimination task (Figure 7.9), with perhaps a small negative effect on levels of best performance reached.

It is interesting to note that the plastic networks with development perform similarly to those without. From this we can deduce that the reason for the poor performance of the plastic networks is not related to the increased length of the transient phase in their dynamics that was identified in Chapter 6. It might be thought that the continual change in network behaviour during the long transient phase contributed to the poor performance by adding noise to the genotype-to-phenotype mapping, but the good performance of non-plastic networks after development shows that the long transient phase is not in itself to blame. It appears from this that it is the presence of active plasticity which is the important factor in causing poor performance on the discrimination task.

7.5.3 Experiment 3: Comparison of pre-conditioned networks with centre-crossing networks

The improved evolvability of the non-plastic networks after a developmental period on the ball-catching task might be because the homeostatic plasticity moves the networks towards the centre-crossing condition during development, as we showed in Chapters 5 and 6. This would tie in with the improved evolvability of centre-crossing networks on a pattern generation task that was demonstrated by Mathayomchan and Beer [127].

Evolutionary runs were performed using the same method as the previous experiments, using non-plastic networks in which the centre-crossing condition enforced under two different regimes. The first regime seeded the initial generation of each genetic algorithm with genotypes calculated to give centre-crossing networks, but then allowed the evolutionary process to move network parameters away from centre-crossing neighbourhood by mutation. Thus in this regime the centre-crossing condition was only enforced in the first generation and was then subject to evolutionary change. The second regime enforced the centre-crossing condition all through the evolutionary run, by calculating the necessary bias terms from the genetically specified weights at every agent initialisation.

Figures 7.10 and 7.11 show the results achieved on the ball-catching and discrimination tasks respectively. The figures show performance in runs with the centre-crossing

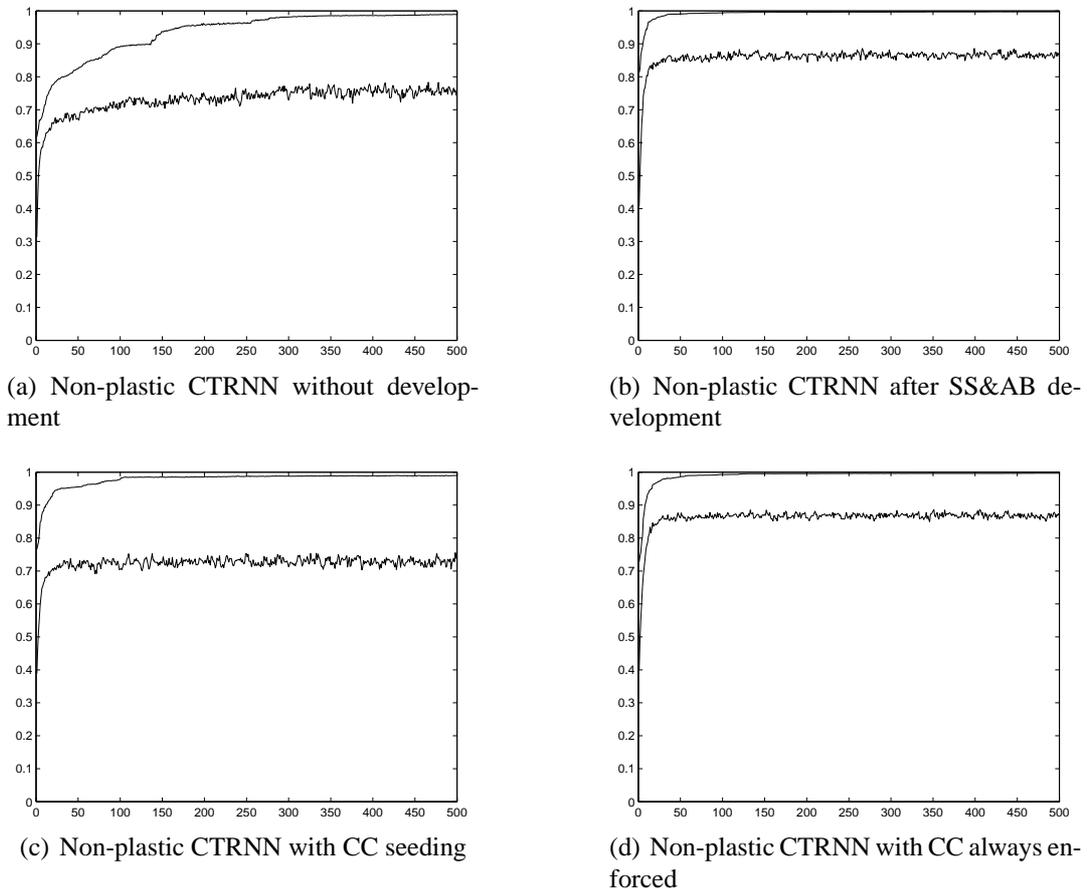


Figure 7.10: *Ball-catching task*. Plots show mean performance over 10 evolutionary runs of non-plastic CTRNNs where the initial population is seeded with centre-crossing networks and also where the centre-crossing condition is enforced for every network in every generation. Also shown is performance of non-plastic CTRNNs with and without homeostatic plastic development for comparison. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. SS - synaptic scaling, AB - adaptive bias, CC - centre crossing. Best performance is shown by centre-crossing networks and by networks conditioned with homeostatic plastic development, which perform to similar levels.

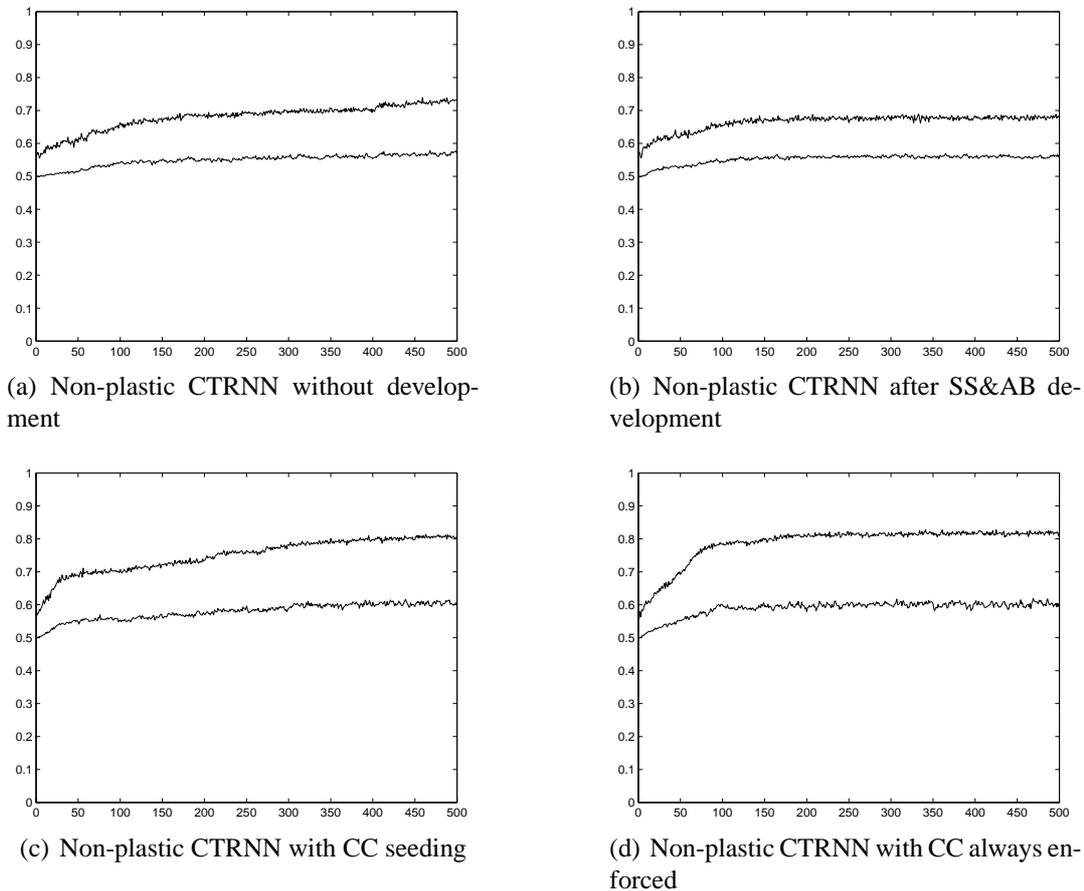


Figure 7.11: *Discrimination task*. Plots show mean performance over 10 evolutionary runs of non-plastic CTRNNs where the initial population is seeded with centre-crossing networks and also where the centre-crossing condition is enforced for every network in every generation. Also shown is performance of non-plastic CTRNNs with and without homeostatic plastic development for comparison. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. SS - synaptic scaling, AB - adaptive bias, CC - centre crossing. Centre-crossing networks perform best on this task, showing better levels of mean and elite fitness than standard CTRNNs and networks with homeostatic plastic development.

condition enforced under the two regimes described above, and also the performance of standard non-plastic CTRNNs and the non-plastic CTRNNs after development (reproduced from Figures 7.2, 7.4, 7.8 and 7.9) for comparison.

In the ball-catching task, centre-crossing networks out-perform the standard networks when enforced in both the regimes described above. The best performance is by networks where the centre-crossing condition is always enforced, which show a similar level of performance to the non-plastic networks that have been conditioned by homeostatic plastic development. This similarity lends support to the hypothesis that the good performance of the non-plastic networks after development is because the homeostatic plastic moves them towards the centre-crossing condition, where they are in a behaviour-rich region of parameter space and poised to behave.

The situation for the discrimination task is a little different. Centre-crossing networks created under both regimes out-perform the other types of network. The best performance is again by the networks where the centre-crossing condition is always enforced.

Taken together the results are a little confusing. In one task, centre-crossing networks and networks where homeostatic plastic development moves them towards the centre-crossing condition show similar levels of success, suggesting that proximity to the centre-crossing condition is the source of their success. In the other task the centre-crossing networks show a significant improvement over the standard networks, whereas the developed networks show a performance decrease.

We know from previous work that homeostatic plasticity moves networks towards the centre-crossing condition. The results for the ball-catching task suggest that this accounts for the success of the pre-conditioned networks. The pre-conditioned networks in the discrimination scenario will still be close to the centre-crossing condition, which suggests that their poor performance compared to centre-crossing networks found by pre-calculation of bias terms must be related to some feature of the developmental process. This finding would appear to contradict our conjecture in the discussion of results from Experiment 2 that the development process itself did not have a negative effect on performance.

7.5.4 Experiment 4: Evolvability of larger non-plastic CTRNNs

One possible explanation for the poor evolvability performance of the homeostatic plastic CTRNNs in some of the previous experiments is that the increased dimensionality of their dynamics makes their behaviour too complex and the intra-network interactions too strong for evolution to be successful.

The complexity of a dynamical system is a difficult property to quantify. To do so in a

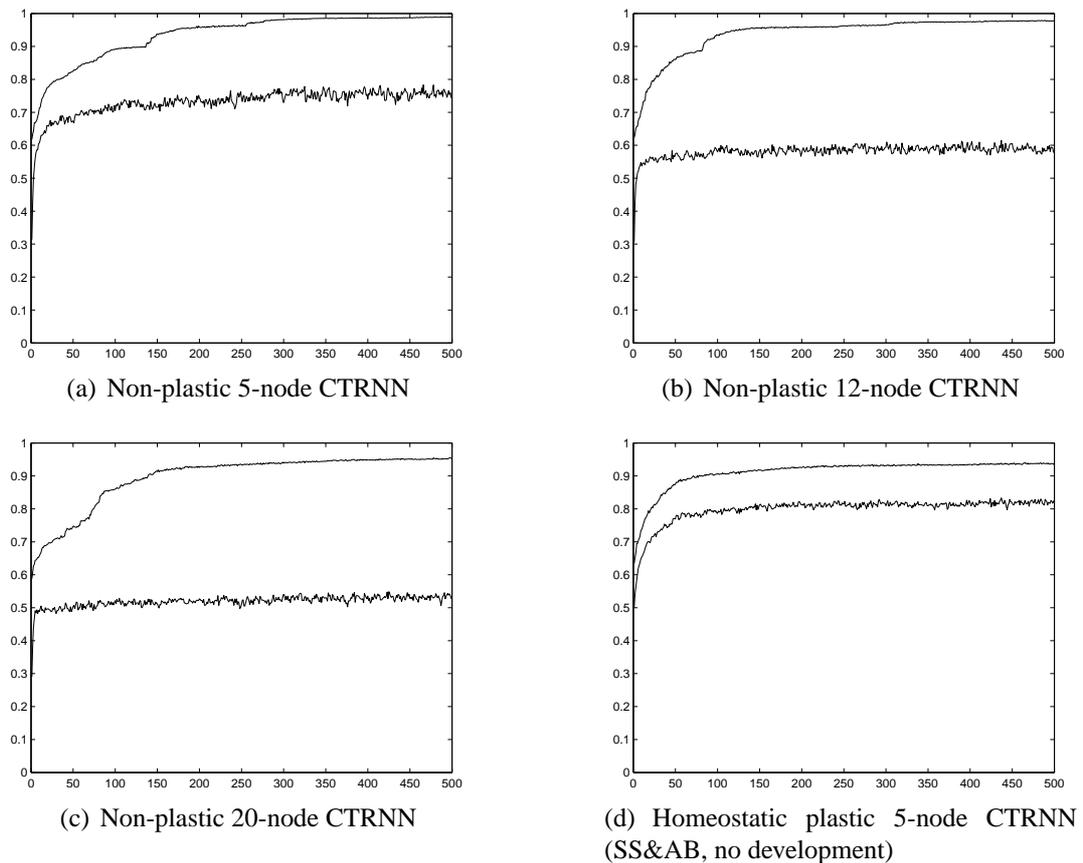


Figure 7.12: *Ball-catching task*. Plots show mean performance over 10 evolutionary runs of non-plastic CTRNNs of different sizes. Also shown is performance of networks with active homeostatic plasticity for comparison. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. SS - synaptic scaling, AB - adaptive bias. Increasing the size of the network decreases evolvability, with lower levels of fitness reached by larger networks.

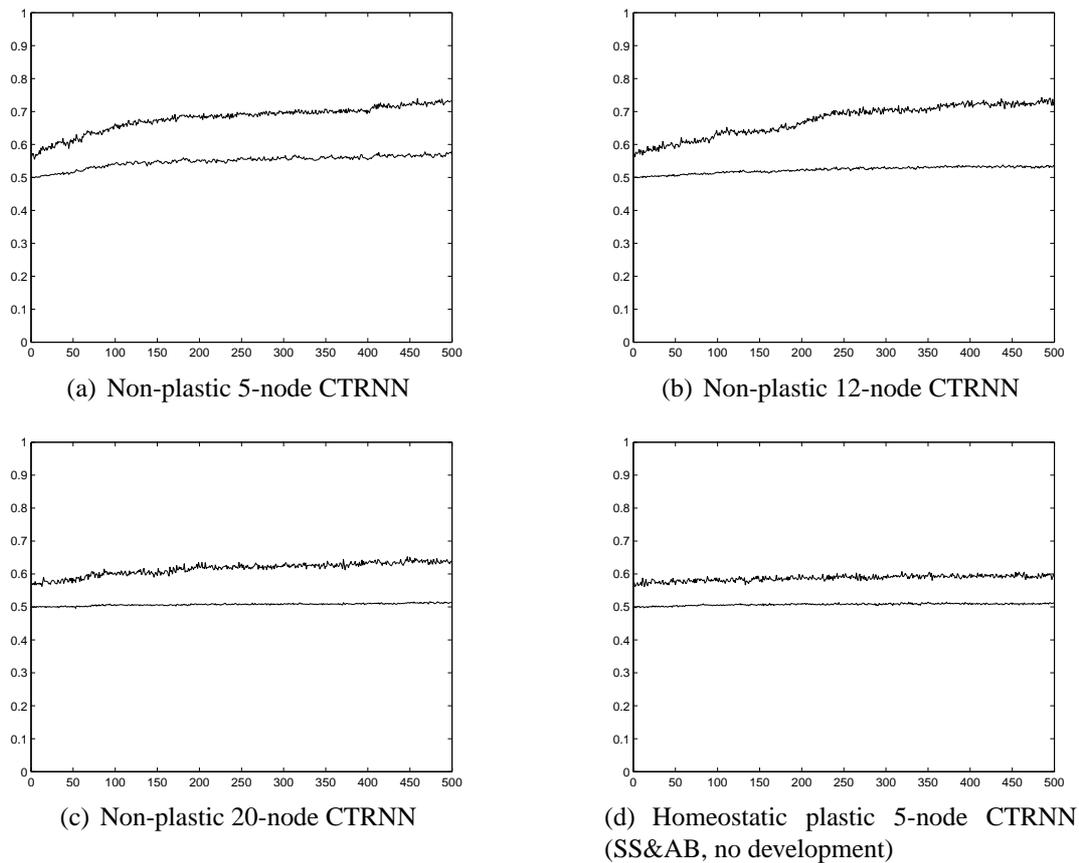


Figure 7.13: *Discrimination task*. Plots show mean performance over 10 evolutionary runs of non-plastic CTRNNs of different sizes. Also shown is performance of networks with active homeostatic plasticity for comparison. In each plot, upper line is fitness score of best individual in population, lower line is mean fitness score for population. Both are plotted against generational time. SS - synaptic scaling, AB - adaptive bias. Increasing network size reduces evolvability, with larger networks reaching lower levels of fitness.

rigorous fashion is not straightforward, but we can perhaps perform some useful informal testing of our hypothesis by testing the evolvability of large non-plastic CTRNNs. By adding more nodes to a network we potentially add more dimensions to its dynamics. Thus we may to some extent test the idea that greater dimensionality in the autonomous dynamics of a system makes it harder to evolve.

Figures 7.12 and 7.13 show the results from evolutionary runs with non-plastic CTRNNs with 12 and 20 nodes. Results for 5-node plastic and non-plastic networks are reproduced from Figures 7.2 and 7.4 for comparison. It can be seen from these plots that increasing the size of the network reduces evolvability, with lower fitness values being reached for both ball-catching and discrimination tasks.

We can note from these runs that adding potential dimensions of variation to the network dynamics in these cases reduces evolvability, and tentatively suggest that a similar effect could occur when dimensions are added by including active plasticity. An N -node non-plastic CTRNN has up to N dimensions in the dynamics, whereas an N -node network with both adaptive bias and synaptic scaling active can have up to $N^2 + 2N$ dimensions. The large difference in the number of degrees of freedom of the plastic networks compared to non-plastic networks may explain the poor performance of the plastic networks.

It should be noted that counting the dimensions in which a system may vary is a very weak method of comparison. The nature of the coupling between the different variables in a system is obviously hugely important in classifying its behaviour. This is another reason why comparing plastic networks to non-plastic networks is not straightforward. Even if they have the same number of dimensions of variation, the interactions between these variables are very different. For this reason we must exercise caution when drawing conclusions about plastic networks from the behaviour of non-plastic networks with greater numbers of nodes.

However, we can observe here that adding dimensions makes evolving good controllers more difficult in some cases and speculate that this may cause the poor performance of the plastic networks.

7.6 Summary

From the results of the experiments described above we can observe that homeostatic plastic CTRNNs reach lower levels of best fitness than non-plastic CTRNNs on both the ball-catching and discrimination tasks, but that the plastic networks are more consistent than the non-plastic networks, with higher population mean fitness and greater consistency between evolutionary runs. The plastic networks also evolve more quickly in the early stages of the evolutionary runs for the ball-catching task, but they always perform

badly on the discrimination task.

The use of a developmental period improves the performance of the homeostatic plastic CTRNNs on both tasks, while the use of a developmental period to pre-condition non-plastic networks gives a large improvement in evolutionary performance on both tasks also. Centre-crossing networks give a similar level of performance to conditioned non-plastic networks on the ball-catching task, but perform better on the discrimination task.

Finally, the evolutionary performance of non-plastic networks decreases as network size increases.

7.7 Implications for robotics

These results deliver a mixed verdict on the utility of homeostatic plasticity for improving evolvability. While homeostatic plastic networks can be more consistent and evolve more quickly to begin with, they are ultimately out-performed by non-plastic networks. It thus appears that the only evolvability benefit that can be claimed by homeostatic plastic CTRNNs is quick early progress and consistency, and even these results are not seen in the discrimination task.

However, when used as a mechanism for pre-conditioning non-plastic networks prior to the commencement of fitness testing, we see an improvement in performance on the ball-catching task, although little difference is seen in the results for the discrimination task. These results seem to suggest a useful role for homeostatic plasticity for improving evolvability. This role is undermined, however, by the greater success of centre-crossing CTRNNs on both tasks.

Centre-crossing networks are the most successful kind of network on both of the evolutionary tasks, and the recommendation from our experiments is that the best way to improve evolutionary progress is to use this kind of network. It seems likely that the good performance of the non-plastic networks after homeostatic plastic development on the ball-catching task is at least in part because the homeostatic plasticity moves the networks to the centre-crossing condition before the trial begins. If there is a simple calculation that can be performed in place of the developmental period with similar effect (i.e., Equation 3.4) then the computational expense of using homeostatic plastic development must count against it.

The most likely explanation for the reduced evolvability of the plastic networks is that their dynamics are simply too complicated for evolutionary search to be able to optimise their parameters. The experiments performed with larger non-plastic CTRNNs showed that as the networks grew larger their evolutionary performance was reduced. We can take this as evidence that more complicated network dynamics reduces evolvability, with

some caveats relating to the qualitatively different nature of plastic networks to non-plastic networks.

One complication is that the genetic algorithm used here allows the possibility of the initial parameters of a plastic network being set by evolution. This could potentially make the plastic mechanisms redundant, since careful parameterisation could prevent homeostasis being lost during normal behaviour. It would be interesting to study the effects of homeostatic plasticity on evolutionary progress in a scenario where network parameters were initialised randomly during each trial, thus ensuring a greater role for the plastic mechanisms in agent performance.

The results achieved do not suggest any immediate evolvability benefit from using homeostatic plasticity in CTRNNs, except when used as a developmental mechanism, in which case an equivalent effect can be achieved with less computational cost by using centre-crossing networks.

However, the experiments we have performed are very limited in scope and should be seen only as the beginnings of an investigation into the evolvability of homeostatic plastic networks. As we discussed above, evolvability testing is not straightforward and more data is required before we can claim a good understanding of the evolvability of homeostatic plastic CTRNNs. It may be that other robotics scenarios, using other kinds of behavioural task or different genetic algorithms, might deliver contradictory results to those seen here. This is likely given the significant differences in the results from the two very similar tasks shown here.

Also, looking for direct evolvability benefits from homeostatic plasticity may be misguided. Homeostatic plasticity is thought by neuroscientists to play a role in regulating network function against perturbations from other types of plastic mechanism during the lifetime of an individual. It may be that homeostatic plasticity in artificial neural networks will be useful performing a similar function in robot brains, maintaining sensitivity in the face of perturbations during the agent lifetime.

These perturbations may be internal, such as changes to network structure caused by other plastic mechanisms. For example, homeostatic plasticity may allow a robot controller to avoid node saturation and maintain network sensitivity in the face of destabilising positive feedbacks caused by Hebbian learning.

Alternatively, perturbations may be environmental. For example, consider the case where a light-sensitive robot must behave in a variety of different environments, each with a different level of ambient light. While sensor nodes in a non-plastic CTRNN might have difficulty in adjusting to the different levels of ambient light, being over-stimulated in bright conditions and under-stimulated in dull conditions, sensor nodes using homeostatic plasticity would be able to adjust their sensitivity to the current conditions and maintain

their functional response.¹

The bottom line is that we simply do not have enough data to draw any major conclusions concerning the utility of homeostatic plasticity for improving evolvability. Here we have presented results from one implementation of homeostatic plasticity, tested on two tasks on a single simulated robot platform. Much more work is needed to develop a fuller understanding. The results presented above do not show much benefit to using homeostatic plasticity, but maybe that is because the situations have not demanded it. If an evolutionary scenario calls for maintenance of network sensitivity and avoidance of saturation in the face of lifetime perturbations, it seems likely that homeostatic plasticity will be of use.

7.8 Implications for neuroscience

The most evolvable networks were the centre-crossing CTRNNs, followed by the non-plastic networks after a developmental period. The developmental period actually creates centre-crossing networks from a non-centre crossing starting position, without the explicit calculation of bias terms that was used for the ‘traditional’ centre-crossing networks used by Mathayomchan and Beer [127].

Biological nerve cells do not engage in abstract mathematics, and are thus unable to create centre-crossing networks by pre-calculation of bias terms, even supposing that a direct biological counterpart to the bias term existed. If centre-crossing networks (that is, sensitive networks consisting of neurons that respond strongly to input) are beneficial, then maybe homeostatic plasticity is a means of creating them that is accessible to biological evolution. Homeostatic plasticity is thought to have a significant role during the development of the mammalian brain [172], and perhaps one of its functions is to move neuronal networks towards a condition where they are sensitive and respond to input, perhaps in the face of pressures from other forms of neural plasticity.

¹The ability of homeostatic plastic CTRNNs to maintain their sensitivity in the face of changing base levels of external input was demonstrated in a published paper based on work in this thesis. See Williams (2004) ‘Homeostatic plasticity in recurrent neural networks’ as referred to in the front matter of this thesis.

Chapter 8

Discussion of Part III

8.1 Overview

In this chapter we will try to summarise and discuss the material presented in Part III, looking at the main findings and their implications for robotics and neuroscience. We will review the methods and approach used, before going on to outline ideas for future work.

8.2 Summary of results

8.2.1 Chapter 3: Background

In Chapter 3 we reviewed the relevant literature concerning homeostatic plasticity both in biological neuronal networks and in artificial neural networks applied to robotics. We introduced evolutionary robotics and continuous-time recurrent neural networks (CTRNNs), that are both fundamental to the work presented in later chapters. We saw that homeostatic plastic mechanisms have recently become an increased focus of research in the neurosciences, and that their use in artificial neural networks for robot control is almost unexplored. Where homeostatic mechanisms have been used with robot controllers, none of the reported work covers the questions we wish to address.

Di Paolo [46,48] and Balaam [8] were interested in Ashby's idea of homeostatic adaptation [7], and accordingly their implementation did not reflect the picture of homeostatic plasticity favoured by neuroscientists.

Another investigation by Di Paolo into the use of homeostatic oscillators [47] is the closest prior study to the current work in terms of mechanism, but was again aimed at the question of behavioural robustness rather than an investigation of the homeostatic mechanisms identified by neuroscientists. While the adaptive mechanism used in [47] is similar to the adaptive bias mechanism used here, the different timescale and research goal mean that there is little overlap with the work presented in this thesis.

Hoinville and Henaff [84] were also interested in robustness, and while the mechanisms they employed were intended to reflect those seen in biological brains, there were significant methodological problems with their implementation and we can learn little from their work.

The centre-crossing networks presented by Mathayomchan and Beer [127] are relevant here, but cannot be said to be homeostatic mechanisms. Although we showed in Chapter 5 that homeostatic plasticity creates centre-crossing networks, the networks used in [127] were created by explicit pre-calculation of parameters and employed no homeostatic mechanisms.

With this background the aims of the current work were identified. The first task was to develop mechanisms to implement homeostatic plasticity in CTRNNs. Subsequently we wanted to see if homeostatic plasticity solved the problem of node saturation, before going on to investigate its effects on CTRNNs at the levels of node, network and agent. Since this work falls mostly under the banner of evolutionary robotics, the final aim was to look at the impact of homeostatic plasticity on evolvability.

8.2.2 Chapter 4: Can homeostatic plasticity be implemented in CTRNNs?

Chapter 4 covered how homeostatic plasticity could be implemented into CTRNNs. The functional effects of the different homeostatic mechanisms seen in biological brains were discussed, and mechanisms acting to scale synaptic weights and alter the intrinsic excitability of neurons were identified as being most suitable for use with CTRNNs. Simple mathematical formulae to govern plastic change were developed, making use of Di Paolo's plastic facilitation mechanism (Equation 4.1) [46], and implementing homeostatic plasticity in CTRNNs in the form of mechanisms for adaptive bias (Equation 4.3) and synaptic scaling (Equation 4.2).

The chapter finished with some discussion of which level of abstraction was appropriate when implementing homeostatic plasticity in CTRNNs, concluding that the level of realism chosen should be (and in fact was) roughly equivalent to the level of abstraction of CTRNNs compared to real brains. The homeostatic plastic networks used here thus retain the same level of biological relevance as CTRNNs.

8.2.3 Chapter 5: Does homeostatic plasticity prevent node saturation?

Chapter 5 opened with some analysis of the nature of the constraint imposed on single-node CTRNNs by specifying a target range for firing rates. It was shown that in most cases, homeostatic plasticity will move the node to a constraint-satisfying region of parameter space, although in some cases there is no stable equilibrium that allows this and nodes may display a slow oscillation as the plastic mechanism continually seeks to satisfy the constraint.

The concept of a constraint-satisfying region of parameter space was generalised to multi-node networks, and it was noted that in a functioning network homeostatic plasticity may move the network through parameter space towards the constraint-satisfying region without ever reaching it, due to the continually changing inputs the network receives. The character of homeostatic constraint satisfaction in a network as an *ongoing process* was discussed.

8.2.4 Chapter 6: What effect does homeostatic plasticity have on network behaviour?

In Chapter 6 we looked at the effect of homeostatic plasticity on CTRNNs at the levels of nodes, networks and agents. This was measured on metrics of sensitivity to input, signal propagation, and likelihood of oscillatory dynamics. The behaviour of a light-sensitive robot controlled by a homeostatic plastic CTRNN was simulated and observed.

Before any experimental work was done it was noted that adding plasticity to a network creates a new class of controller, and is not simply an additive change to the existing network. This raises difficulties in fair comparison, and the method of adiabatic approximation for comparing the behaviour of non-plastic networks with plastic networks was discussed and used throughout the experiments.

It was shown that homeostatic plasticity allows slow oscillations to occur in a single node, and causes individual nodes to become more sensitive to external input by moving them closer to the centre-crossing condition.

In networks, homeostatic plasticity was shown to add a slower timescale to network dynamics and cause a large increase in the time taken for network dynamics to converge to an attractor from a random initial state. Signal propagation was improved in fully connected, randomly connected, and feedforward CTRNN architectures. Oscillations were much more likely to occur in both homeostatic plastic networks and in non-plastic networks after a period of homeostatic plastic development. Oscillations in firing rates were

made more likely by homeostatic plastic development because they made networks more excitable and allowed signals to propagate further. Plastic networks behaved similarly, but also had the possibility of slower oscillations involving changes in weights and biases, caused by individual nodes failing to satisfy the homeostatic constraint.

Homeostatic plasticity made the interaction between the photo-sensitive agent and a light source in its environment much stronger both when the plasticity was left active and also when used as a pre-conditioning mechanism for the agent controller. This was demonstrated by motion plots of agent's controlled by randomly parameterised networks interacting with a series of light sources. Standard CTRNN controllers typically ignored the light sources, while plastic controllers showed obvious and continuous interaction.

8.2.5 Chapter 7: Does homeostatic plasticity make CTRNNs more evolvable?

Chapter 7 opened with an argument for why homeostatic plastic networks might be more evolvable than non-plastic networks, citing their properties of increased sensitivity, better signal propagation and increased likelihood of oscillations. These properties are thought to be important in a number of biological motor control systems and could therefore be useful for artificial agents.

A series of evolvability experiments was then reported, after some preliminary discussion of the difficulties in measuring evolvability caused by the contingent nature of evolution. These experiments used a simulated agent whose control network was evolved with a genetic algorithm to catch and discriminate between different types of falling shape.

The results of the evolvability experiments were not conclusive. In terms of best fitness, networks with active homeostatic plasticity generally performed worse than non-plastic networks. Although they showed quicker progress early in evolutionary runs and greater consistency in reaching a reasonable level of performance, they were eventually out-performed by non-plastic networks.

When homeostatic plasticity was used just as a developmental mechanism for the robot control networks, that is, when it was applied for a period prior to the fitness trial and then switched off again, it had significant beneficial effects on one of the tasks and slight beneficial effects on the other task. The developmental period caused quicker evolution of good controllers.

It was conjectured that this benefit might be because the developmental process moved the control network to the centre-crossing condition, thus leaving it poised to behave. To test this hypothesis, some runs were performed with centre-crossing networks created by calculating bias terms as in the original work by Mathayomchan and Beer [127]. These

centre-crossing networks proved to be the most successful type of controller on both the tasks. They were better than the plastic and non-plastic networks in terms of speed of evolution and best level of performance reached. They also achieved equivalent or better performance than the non-plastic networks conditioned by a developmental period of homeostatic plasticity. Since the centre-crossing networks achieved at least equivalent performance to the developed networks, without incurring the computational cost of the developmental process, the centre-crossing networks were held to be more evolvable than the developed networks on the tasks investigated here.

Chapter 7 concluded with some discussion of the significance of the results presented. It was pointed out that due to time and space constraints the question of evolvability had not been fully addressed, and that more work was needed to gain a more complete understanding. Also, since neuroscientists believe that the function of homeostatic plasticity in biological nervous systems is to control the destabilising effects of other plastic mechanisms, it was suggested that homeostatic plasticity might be most useful playing a similar role in artificial neural networks. If the evolutionary task involved disruptions to network function by alternative forms of plasticity or from environmental perturbation, homeostatic plastic networks may be more evolvable and perform more strongly than they did on the static control tasks employed here.

8.3 Implications for robotics

From the perspective of robotics, we are interested in whether or not we have learned anything from our studies of homeostatic plasticity that can be used to help us build better robots more quickly. The results we have presented suggest a number of ways in which homeostatic plastic CTRNNs might be useful.

Homeostatic plasticity as implemented above operates on a slower timescale than neural activation dynamics. This may be useful in situations where the environment or task faced by a robot requires action on a wider range of timescales than those easily allowed by standard CTRNNs. While standard CTRNNs can implement dynamics on timescales slower than that of individual nodes by carefully designed network structure, homeostatic plasticity may allow easier access to a wider range of timescales using fewer nodes.

Also, the slow oscillations that can be displayed by single homeostatic plastic nodes may allow some useful form of switching. The flip-flop behaviour of these nodes as they alternate between high and low activations may be useful as a form of binary switching that could perhaps be used to trigger different network behaviours.

The network-level properties resulting from the application of homeostatic plasticity (i.e., increased sensitivity, better signal propagation and more oscillations) seem likely

to be useful for a variety of robot control tasks. However, since these properties are also demonstrated by centre-crossing networks and are thus available without the computational overhead of homeostatic plastic development, it seems likely that homeostatic plasticity will only be useful in situations where these properties must be maintained in the face of perturbation. This is supported by the evolvability results shown in Chapter 7, where centre-crossing networks were the most successful controller on the static tasks used.

It is easy to imagine situations where the ability to maintain sensitivity and signal propagation in a network in the face of perturbations will be useful. One obvious case is networks where other forms of plasticity are active. For example, if Hebbian learning rules are altering synaptic weights in response to patterns of associated activity, this can cause positive feedback and runaway weight change. This would lead to extreme weight values and network saturation if left to run unchecked. Homeostatic plasticity could provide a mechanism by which such changes were counter-balanced. This balance would be dynamically maintained online during the agent lifetime, keeping networks sensitive and ensuring some form of interaction between the agent and its environment.

A less obvious situation where the online maintenance of sensitivity and a non-trivial coupling to the environment would be useful is the transfer of controllers from simulation to hardware. As mentioned briefly in the introduction to evolutionary robotics in Chapter 3, it can often be difficult to successfully transplant robot controllers that have been evolved in simulation to the hardware they are supposedly designed to control. This problem of bridging the ‘reality gap’ [89] might be ameliorated by the inclusion of homeostatic plasticity, since the plastic mechanisms should allow some of the inevitable differences between the simulated hardware and the actuality to be overcome.¹ For instance, if the level of gain in the interface between controller and hardware is not the same as the amount in the simulation, the homeostatic plasticity should allow the controller to adjust the sensitivity of the relevant nodes to compensate. A similar situation might occur in cases where the real environment does not provide the same feedback as its simulated counterpart. If the levels of stimulation received by the robot sensors are consistently too high or too low, then homeostatic plasticity in the sensor nodes should allow for some compensation to overcome this.

The ability to adjust excitability of sensor nodes should also be useful where the task

¹For instance, Thompson [165] highlights unforeseen problems related to temperature when transferring artificially evolved field-programmable gate array (FPGA) circuits from simulation to reality. The temperatures generated during the operation of the FPGA in some cases caused circuit malfunctions; different ambient environmental temperatures also caused negative effects. Interestingly, Thompson proposed solutions based on biological schemes of thermo-regulation.

faced by the robot involves operating in different environments.² For instance, if the photo-sensitive robot used in Chapter 6 was placed into a very bright or very dim environment, it may fail to detect relevant light gradients due to its sensors being consistently under- or over-stimulated by the extreme levels of ambient light. If the sensor nodes incorporated homeostatic plasticity they might be able to adjust their baseline excitability to match the environment, continuing to distinguish different light intensities and give useful information to the network.

The results we have presented in this thesis show that homeostatic plasticity can be implemented in CTRNNs, and explore some of the effects of doing so. While the inclusion of active homeostatic plasticity makes networks less evolvable in static (internal and external) environments, there is a strong argument in favour of some utility for homeostatic plasticity in situations where network functionality must be maintained in the face of perturbations.

8.4 Implications for neuroscience

While pure engineering roboticists are not directly interested in the biological relevance of their robot controllers, except in so far as it can help them to design better robots, the enterprises of biomimetics, neuroinformatics and cognitive science are directly concerned with what we can learn about biology through the construction of this kind of artificial simulacrum. It is important to be clear about the limitations of the work presented here in this respect.

The first thing to state when considering the implications of the work presented in this thesis for neuroscience is that the models of neuronal networks and homeostatic mechanisms used here are highly abstracted. The details and complexity of biological brains have been sacrificed in favour of analytical and computational tractability. CTRNNs are not brains, genetic algorithms are not biological evolution.

However, in the sense that they are distributed network processors of rate-based information, where individual nodes fire at a rate determined as a sigmoidal function of multiple inputs, and where activation in the network is persistent over time, CTRNNs are reasonable, if simple, models of brains. We can follow Beer's analogy of the 'frictionless brain' [11] and seek to gain insights from the functioning of these artificial networks that may guide experiments by 'real' neuroscientists. For this purpose we can make several observations about the results from our experiments with homeostatic plastic CTRNNs.

The first result to note is that while homeostatic plasticity regulates the activity of

²See again Thompson [165].

nodes to a moderate level in the long term, the result of this is greater response to input. Homeostatic plasticity regulates the activity of individual nodes so that they are not saturated, and this means that their inputs have a larger effect on their firing output. This leads to improved signal propagation and more oscillations at the network level, which can arguably be seen as the converse to the regulation of network activity. This seems to go against the view in neuroscience that homeostatic plasticity stabilises network activity.

Another prevalent idea in neuroscience is that homeostatic plasticity acts as a counterbalance to the destabilising effects of Hebbian learning. However, it seems that Hebbian learning in CTRNNs actually has a stabilising, not destabilising, effect. When Hebbian rules are applied to CTRNNs, a common result is positive feedback between activity and weight change, so that weights are quickly forced to extreme values. This results in node saturation and a loss of response from individual nodes, leading to a general lack of signal propagation at the network level. This can be viewed as a form of stabilisation of network dynamics, albeit one which is likely to cause a loss of functionality.

So we have the neuroscientific opinion that homeostatic plasticity stabilises network activity and Hebbian plasticity destabilises network activity, compared to our findings that the *converse* is what actually occurs in CTRNNs. Homeostatic plasticity in CTRNNs potentially destabilises network activity by allowing signals to propagate further and making oscillations more likely, while Hebbian rules often lead to the desensitising and thereby stabilising effect of node saturation.

This conflict suggests that it would be useful for neuroscientists to study the hierarchical effects of both Hebbian and homeostatic plastic mechanisms to try and elucidate whether or not the network-level effects of node-level Hebbian and homeostatic plastic mechanisms are what they are purported to be. Such experiments are likely to be difficult because of the inherent complexity of biological neuronal networks, and simulation modelling at a higher level of biological accuracy may be the most practical initial route to the required data. The results from the simple simulations we have undertaken here suggest that this may be a fruitful line of enquiry.

It appears from our results that homeostatic plasticity and Hebbian learning act in opposite ways. In view of this, we may be able to identify a higher level of homeostatic regulation at the network level, where the regulated quantity is the level of activity in the network, rather than of individual nodes. Hebbian learning causes node saturation, which reduces the amount of activity in the network. Homeostatic plasticity increases the response of individual nodes and raises the level of network activity. Since homeostatic plasticity raises network activity and Hebbian learning reduces network activity, it seems likely that the two types of mechanism will balance each other out, providing a measure of long term stability in the overall level of network activity. In effect, Hebbian learning and

node-level homeostatic plasticity may combine to maintain network-level homeostasis of activity.

8.5 Review and Future Work

It would have been advantageous to have run more experiments on the evolvability of homeostatic plastic CTRNNs, since the data set generated in Chapter 7 is limited. While it can be argued that there is rarely ‘enough’ data in any experiment, that more data is always useful and that the law of diminishing returns applies, two evolutionary scenarios are not enough to gain a true picture of the evolvability of homeostatic plastic CTRNNs. In particular, it would have been very interesting to see how homeostatic plastic networks performed in situations where there was a need for the online maintenance of network function in the face of perturbations.

This leads to the other obvious omission from the thesis, which is the lack of any experimental exploration of the interaction between homeostatic plasticity and Hebbian learning. While reference has been made to this interaction in the text, no study of Hebbian learning in CTRNNs has been undertaken here. The simple reason for this is that such a study is a whole topic in itself, requiring preliminary studies of the literature, development of mechanisms, measurements of the effects of Hebbian rules on CTRNN dynamics, and examination of the utility of Hebbian rules in evolutionary scenarios. This is an equivalent body of work to the entirety of Part III, and would need to be performed before any study of the interaction between homeostatic plasticity and Hebbian rules could reasonably be undertaken. Such a study is beyond the scope of this thesis, and while it may form the topic of future work, the reader will for now have to be content with the passing discussion presented here.

A weakness in the presented analysis of homeostatic plasticity is that no work was done to look at the nature of evolved plastic controllers. It would have been useful to analyse a number of evolved controllers and find out how the plastic mechanisms contributed to agent behaviour. For instance, is the plasticity used just as a regulatory mechanism or does it form an integral part of the dynamics of the evolved controllers?

Another area where more work could usefully be done is some study of the effects of homeostatic plasticity on the robustness of evolved solutions. This is related to the earlier conjecture about the utility of homeostatic plasticity in allowing networks to overcome perturbations, and well-designed experiments could perhaps attack both questions simultaneously. Results concerning behavioural robustness would allow the results given here to link to the existing work of, e.g., Di Paolo [46–48, 50], Balaam [8], and Hoinville & Henaff [84].

Homeostatic plasticity clearly exists in biological brains and must therefore offer some adaptive benefit that caused its evolution. Our experiments so far have not shown a positive relationship between the presence of homeostatic plasticity and evolvability. Further experiments may demonstrate such a relationship, but our conclusion from current results is that the benefit of homeostatic plasticity must lie elsewhere.

Finally, we should review the overall thrust of the research presented here. Were the right questions addressed?

In Chapter 3 we saw that the literature on homeostatic plasticity in artificial neural networks is small, and that the existing work has largely been aimed at questions of robustness. Thus we feel justified in attempting a more general exploration of how homeostatic plasticity changes network dynamics, and what its functional effects might be. We feel that we have addressed many of the basic questions (e.g., satisfaction of the homeostatic constraint, effect on node/network/agent behaviour), and partially addressed the issue of evolvability. However, there are clearly many functional questions remaining to be answered and we view the work presented here as a first step towards a fuller understanding.

8.6 Conclusions

In our exploratory studies we have shown that homeostatic plasticity can be implemented in CTRNNs. It drives networks towards a region of parameter-space where the homeostatic constraint is satisfied, but it is important to note that it may never reach a stable equilibrium inside this region. Sometimes initialisation dictates that the network falls into a dynamic equilibrium that continually overshoots this region, a limit cycle equilibrium where the constraint is not satisfied. Also, in any functioning network (i.e., one connected to some sensorimotor apparatus) the continual influence of external input will in most cases mean that the action of homeostatic plasticity is best viewed as a *continuously perturbed process* rather than a ballistic trajectory towards some goal state.

Homeostatic plasticity increases node and network sensitivity, improves signal propagation in networks, and makes oscillatory dynamics more likely. It adds a slower timescale to network dynamics and may allow for an approximation to binary switching in networks as a result of slow single-node oscillations. It leaves networks poised and behaviour-rich, increases the strength of interaction between agent and environment, and thus ensures a non-trivial agent-environment coupling.

In the situations tested here, namely ball-catching and discrimination tasks with a simple agent, active homeostatic plasticity did not improve evolvability. Both these tasks involve an unperturbed static environment, and it is possible that homeostatic plasticity

will improve evolvability in situations where perturbation is likely or adaptation to different levels of stimulation is required. Homeostatic plasticity may also have an important role in counter-balancing the positive feedback effects of Hebbian learning.

Our studies also challenge the existing view in neuroscience, that homeostatic plasticity acts to stabilise network dynamics. While it does so at the level of individual nodes, the network-level result of this is increased excitability and greater levels of activity.

Part IV

Homeostatic Regulation of the Environment by the Biota

Chapter 9

Background to Part IV

9.1 Overview

The idea that feedback from the environment shapes the way in which organisms adapt is well established in evolutionary theory. Increasing emphasis is being placed on coevolutionary change, where multiple species adapt in response to feedback from each other. This move away from the idea of a fixed evolutionary environment, and a static fitness landscape for an evolving species, echoes Darwin's original view of evolutionary change as resulting from the interactions between the many species inhabiting the 'entangled bank' [36].

Despite the increasing emphasis on coevolutionary adaptation, the idea that the environment can be altered by the actions of living things (and that coevolution can take place between species and environment, as well as between species and species) has until recently been largely ignored. But with the establishment in evolutionary theory of the ideas of niche construction [102, 103] and the 'extended phenotype' [40], feedback between the environment and the biota is now seen as a bi-directional process.

This shift in evolutionary thinking has been accompanied by the development of a more daring theory concerning the interaction between life and its environment: Gaia theory [124]. Gaia theory claims not only that the feedback between life and the environment works in both directions, but also that life alters the environment to maintain conditions suitable for life. Initially outrageous and still controversial, Gaia theory has slowly gained support from mainstream science to the point where many of its claims are

now well supported.

Some examples of proposed Gaian phenomena include the regulation of local climate by marine algae that influence the formation of clouds over the oceans [5, 27, 114, 122], global temperature regulation by biotic enhancement of rock weathering [152, 153], the maintenance of constant marine salinity and nitrogen:phosphorous ratios by the aquatic biota [110, 140, 176], and efficient nutrient recycling loops created by the collective metabolic processes of microorganisms [176, 187].

Such phenomena are measurable and testable by scientific experiment, and thus provide firm evidence that regulatory loops do exist in nature. But do they point to a general occurrence of regulation in the biosphere? There remain many areas of contention and much research to be done. The mechanisms and processes by which Gaia operates are elusive and are not fully understood. Perhaps the most significant theoretical challenge lies in achieving a reconciliation between Gaia theory and evolutionary biology. Gaia has been criticised by evolutionary theorists because it seems to imply unfavoured mechanisms such as teleology, altruism, or group selection, but Gaians argue that these criticisms are not justified. It is in this area that the work presented in subsequent chapters lies.

Here we examine the hypothesis that the evolving biota regulate their environment to conditions suitable for life. Starting from the view that species coevolve with each other and with their environment, we present two different simulation models of environmental regulation that explore the implications of various hypotheses and assumptions. One model is an extension of the Daisyworld model initially presented by Watson & Lovelock [183], while another model is a piece of original work intended to relax some of the assumptions inherent in the Daisyworld formulation.

This study adds a new dimension to our over-arching theme of homeostatic adaptation, in that the adaptive mechanism at the local level is now evolutionary change by Darwinian mechanisms. The inclusion of evolution as a possible mechanism for homeostatic adaptation is a novel exploration that complements the previous studies, especially as environmental regulation by the biota involves the maintenance of a macro-level homeostasis by micro-level adaptation.

This chapter briefly reviews the literature on Gaia theory, tracing its historical origins before looking at its current status and existing theoretical models. Chapter 10 presents an extended Daisyworld model and some results generated from it. Chapter 11 presents a sketch of the new and original Flask model, which is based on a different set of assumptions to Daisyworld. This sketch is followed in Chapter 12 by some preliminary results from Flask. Part IV concludes in Chapter 13 with some discussion of the results achieved in earlier chapters and conclusions concerning the research questions addressed.

9.2 Gaia theory

This review of Gaia theory will not attempt to exhaustively cover what is a large and broad literature in its entirety. Instead it will give a concise synopsis of the most significant and relevant pieces of work for the current purpose.

9.2.1 Precursors: Vernadsky

The Gaia theory was first named and expounded by Lovelock & Margulis in 1973 [124], but it owes much to the earlier work of the Russian scientist Vladimir Ivanovich Vernadsky (1863-1945). Vernadsky looked beyond traditional views of biology and developed a view of life centred on the *biosphere* [174, 175]. The term ‘biosphere’ had been proposed by geologist Eduard Suess in 1875 to describe the region of the Earth’s surface – air, land, surface rocks and water – in which life existed, differentiating it from the other geological zones of lithosphere, hydrosphere and atmosphere. Vernadsky realised that life existed in parts of all of the other three zones and that the biosphere was not *distinct* from them but rather *included* them. He re-defined the biosphere as the interacting system of life together with its physical environment.

Vernadsky’s concept of the biosphere as a complex interacting system of life and the physical environment was the first recognised theory to give life a major role in the development of the physical environment. He saw living matter as the greatest of geological forces, involved in transporting and transforming matter across oceans and continents, both in the form of living things and as a result of their actions. Vernadsky’s deliberately vague definition of life as geological force was a departure from the Cartesian duality of spirit and matter, life and non-life, and emphasised the notion of life as a process, heavily involved in shaping the atmosphere, hydrosphere and lithosphere.

9.2.2 Early development

In the early 1960s, James Lovelock was working for NASA on methods for the detection of life on other planets. He realised that this could be done by testing the composition of the atmosphere. Lovelock thought that any type of life must take in some form of energy and matter and excrete waste products, and that this process must be mediated by atmospheric transport of these materials. Thus a reliable signature of life on a planet would be an atmosphere in chemical disequilibrium, since the only way in which this could occur would be if some process (i.e., life) was continually pumping mutually reactive chemical products into the atmosphere. On a planet without life, the atmosphere would soon fall to equilibrium as those chemicals capable of reacting did so. Considering the Earth, Love-

lock noted that the chemical composition of the atmosphere had been held stable away from equilibrium for long periods during Earth history and reasoned that this must be due to the cumulative effects of all living things.

Lovelock's hypothesis that the biota regulated the atmosphere to conditions suitable for life was given the name Gaia by a friend (the novelist William Golding) after the Greek earth goddess, and was first published in a 1973 paper co-authored by Lynn Margulis¹ [124]. The Gaia hypothesis attracted criticism, but was refined throughout the 1970s before being published in extended form as the 1979 book 'Gaia' [116].

Lovelock defined Gaia as:

...a complex entity involving the Earth's biosphere, atmosphere, oceans, and soil; the totality constituting a feedback or cybernetic system which seeks an optimal physical and chemical environment for life on this planet. The maintenance of relatively constant conditions by active control may be conveniently described by the term 'homeostasis'. ([116], p.10)

This definition, together with the title of the initial 1973 paper, "Atmospheric homeostasis by and for the biosphere: the Gaia hypothesis" (p.1, [124]), together give a good view of Lovelock's early thinking. He thought that life regulates the atmosphere to provide conditions optimal for life.

9.2.3 Criticisms

Lovelock outlined a number of candidate global regulatory mechanisms in his 1979 book [116], which centred on the constancy of the surface temperature of the Earth (in the face of increasing solar luminosity), the constancy of the composition of the atmosphere (held stable away from equilibrium), and the constancy of the salinity of the oceans. However, it was on theoretical grounds that the Gaia hypothesis attracted most criticism.

Before we look at these criticisms, it is worth first of all noting that by choosing such an emotive name for the Gaia hypothesis, Lovelock may have done the underlying science a disservice. Naming the theory for a Greek earth goddess brings with it connotations of mysticism and hippy sentimentality that may have offended the purist analytical tendencies of the orthodox scientific community. Lovelock did not help to allay these fears by talking of Gaia as a quasi-living superorganism and using descriptive language in papers on Gaia that appeared to allude to sentience. Gaia was adopted as a talisman by various New Age and environmentalist groups, amplifying the doubts felt by hard-headed objective scientists.

¹Margulis was already no stranger to controversial theories, cf., her theory of serial endosymbiosis.

It was in this climate that Dawkins [40] and Doolittle [52] independently made their early scientific criticisms of the Gaia hypothesis. These centred on accusations of teleology, namely that regulation of the biosphere by the biota would require foresight and planning. In any case, contributing to global regulation would involve altruistic sacrifice by participating organisms that would result in them being out-competed by non-contributing ‘cheaters’ and regulation would thus break down. Other criticisms disputed the possibility that feedback mechanisms would be sufficient to create global regulation, argued that such regulatory feedback loops could not have evolved, and claimed that the Gaia hypothesis was untestable and thus non-scientific.

9.2.4 Daisyworld

In response to these criticisms, Watson & Lovelock presented the Daisyworld model [183]. In Daisyworld, populations of black and white daisies compete for space on an artificial planet, with the proportion of the surface area covered in black or white daisies determining the planet's albedo and thus its temperature. It was demonstrated that competition between the two daisy species led to regulation of the temperature of the planet around the optimal level for daisy growth.

Daisyworld answered several of the criticisms aimed at Gaia theory. It demonstrated that global regulation could emerge by a process of positive and negative feedbacks without any foresight or planning, and thus rebutted the teleology criticism. However, since there was no cost to daisy pigmentation, there was no scope for either altruism or selfishness, and the original Daisyworld model could not answer the cheater criticism. Also, the limited amount of genetic variation (daisies were all identical except for pigmentation and could only be black or white) and lack of mutation meant that Daisyworld in its original form did not shed much light on the compatibility of Gaia theory with neo-Darwinian evolutionary biology.

We will cover the Daisyworld model in more detail below.

9.2.5 Multiple Gaia hypotheses

James Kirchner [97] argued that the Gaia debate needed clarification, and that there was not one but several Gaia hypotheses. Kirchner's typology of Gaia hypotheses is listed below in order of increasing strength:

- *Influential Gaia*. The biota has a substantial influence over the aspects of the composition of the abiotic world.

- *Coevolutionary Gaia.* The biota influences the environment and the environment influences the evolution of the biota.
- *Homeostatic Gaia.* The biota influences the abiotic world in a way that is stabilising; the major linkages between the biota and the abiotic world are negative feedback loops.
- *Teleological Gaia.* The atmosphere is kept in homeostasis not just *by* the biota, but *for* the biota; some sense of purpose is implied.
- *Optimizing Gaia.* The biota manipulates the environment to create favourable or optimal conditions for itself.

Kirchner was able to supply quotations and references in support of each of these variant hypotheses from Lovelock's own published work, and pointed out that much of the debate over Gaia might stem from different apprehensions of what the Gaia hypothesis actually stood for. This fragmentation of Gaia theory into many different hypotheses is still much in evidence today. A recent special issue of the *Climatic Change* journal on Gaia theory includes papers by several leading authorities, each of whom independently start their discussion by setting out several competing versions of the Gaia hypothesis [149].

As the Gaia theory has changed over the years, the various hypotheses that have been put forward can be grouped according to whether they are 'weak' or 'strong' [126]. Weak Gaia says that feedback exists in both directions between life and the environment, and that the whole Earth system co-evolves, a view which is now widely accepted. Strong Gaia says that the planet and its lifeforms constitute a single living system that regulates itself to optimal conditions for life. It is this strong Gaian view that attracts most criticism from neo-Darwinians, who say that a unified planetary system cannot have evolved except through competition with other planetary systems.

It seems likely that the truth will lie somewhere between these two extremes. Most people now agree that the biota influences the abiotic world, but few believe it does so purposively; much of the ongoing debate concerns how strong a Gaia theory can be supported by scientific evidence.

9.2.6 Gaia vs Evolutionary biology

Gaia has suffered many criticisms from neo-Darwinian evolutionary biologists over the years, who are mainly concerned with how global regulatory feedback loops could arise by natural selection.

9.2.6.1 Early criticisms

The first complaint against the Gaia hypothesis was that it seemed to imply teleology. However, in 1988 the American Geophysical Union hosted a conference devoted entirely to Gaia [151] and at the meeting Lovelock presented a revised form of the Gaia hypothesis [120] that made particular efforts to refute the claims of teleology. Lovelock's revised hypothesis claimed that Gaia regulated the biosphere by homeostatic feedback loops, and was supported by results from Daisyworld. The teleology criticism appears to have been dropped by the scientific community after this conference, but other theoretical problems remained.

Sophisticated regulatory mechanisms are typically the result of competition and selection between individuals; better regulators make better survivors and better reproducers, and thus regulation is selected for. However, there is no population of life-bearing planets for selection to operate on and no concept of planetary reproduction, so global regulation must result from cooperation and coordination between individuals at a lower level. On the other hand, organisms altruistically contributing to global regulation runs counter to the accepted idea of selfish natural selection; such foolishly generous organisms would be out-competed by non-contributing 'cheaters' who reaped the benefits of environmental regulation without incurring the cost of participation. Gaia theory seemed to imply group selection or costly altruism, and was thus incompatible with modern evolutionary theory.

9.2.6.2 The case for the evolution of regulation

The debate over whether Gaia theory was compatible with neo-Darwinian evolutionary biology continued through the 1990s. Lenton [105] reviewed the issues in the debate, stating the main question to be answered as "...how can self-regulation at the planetary level emerge from natural selection at the individual level?" (p.439, [105]).

Lenton constructed a logical argument for why global regulation must have evolved. Life on Earth has a long history, which might be explained either by the presence of global self-regulating mechanisms (Gaia) or by random chance (anti-Gaia), with an intermediate position accepting that regulatory feedbacks exist but considering their genesis as a matter of good luck rather than an inevitable outcome of evolution. Lenton favoured the Gaian view, pointing to the low probability of life persisting and maintaining a favourable climate in the face of perturbations from planetesimal impacts and volcanic eruptions if there were no regulatory mechanisms.

Lenton then used Daisyworld models to demonstrate that global regulation could emerge from selfish adaptation at the individual level. By introducing mutation of daisy albedo into the Daisyworld scheme he showed that the evolution of stable collective reg-

ulation could occur, but noted that mutation of the preferred growth temperature would lead to evolutionary change that destroyed regulation.

Lenton also argued that the cheater criticism could be answered by noting that many organisms alter their environment not for purposes of planetary regulation but because it brings them some local benefit, and that global regulation could emerge as a byproduct of many species acting in this way; thus there is no conflict between contributing to regulation and acting selfishly, since the two actions are the same.

9.2.6.3 The case *against* the evolution of regulation

A good synopsis of the current disagreements between Gaia theory and evolutionary biology is given by Kirchner [98]. Kirchner is critical of the homeostatic Gaia theory, stating that beyond the accepted view that there is feedback coupling between organisms and the environment, there is little evidence in support of further Gaian claims. He identifies three central propositions to the modern Gaia hypothesis: (1) biological feedbacks to the environment contribute to global homeostasis, (2) biological feedbacks make the environment more suitable for life and (3) biological feedbacks will evolve by natural selection.

Kirchner refutes each of these propositions in turn. He claims that biological feedbacks are not inherently homeostatic and that the feedback between organisms and their environment is just as likely to be positive as negative. Positive feedbacks would be destabilising and would lead to non-homeostatic (anti-Gaian) amplification of changes in climate.

Kirchner then argues that rather than biological feedbacks making the environment more suitable for life, it is much more likely that life simply adapts to its environment, which is in turn altered by the effects of life. If the environment, created by the actions of the biota, appears to be beneficial for life, it is simply that the currently existing biota are well-adapted to the environmental conditions they create. The fit between biota and environment is thus explained by biological adaptation alone and no Gaian mechanisms for 'improving' the environment need be invoked.

Finally, Kirchner argues that although it is possible for Gaian feedbacks to evolve, it is also possible for anti-Gaian feedbacks to evolve. He casts doubt on the Daisyworld model as a proof of a tendency for homeostatic feedbacks to evolve, pointing out that the regulation in Daisyworld relies on the assumption that a beneficial adaptation for an individual daisy will by default have a regulatory effect on the environment. The situation in the real world is not the same, since evolution will favour any trait that offers a reproductive advantage to carriers over non-carriers. The effect of such a trait on the environment is largely irrelevant since this will affect both carriers and non-carriers in the

same way. This means that Gaian and anti-Gaian feedbacks are equally likely to evolve.

9.2.6.4 By-product Gaia

The ‘by-product Gaia’ version of the theory developed by Volk [176] avoided many of the criticisms from evolutionary biology by positing a weaker hypothesis based on life’s involvement in nutrient recycling in the biosphere.

Volk took the Gaian idea that ‘life begets life’ and developed an alternative viewpoint by arguing that the key Gaian phenomenon was the increased availability of nutrients in the environment that resulted from biotic influences. In an abiotic environment, Volk argued, nutrients were only available as a result of chemical and physical processes such as weathering and volcanic out-gassing. In a world with life, nutrients were recycled many times via metabolic processes and their availability was greatly increased. Life did indeed beget life, by increasing the food supply in the environment compared with the lifeless alternative.

Volk felt that any homeostasis of the environment was a fortuitous result of nutrient recycling and his view focused on metabolic by-products rather than regulation. Thus Volk’s view of Gaia was not vulnerable to the same criticisms as the original ‘homeostatic’ Gaia. There was no altruistic cooperative regulation to be vulnerable to ‘cheaters’, and feedback loops involving metabolism and biota-enhanced physical processes were fully compatible with selfish natural selection.

9.2.7 Gaia - an organism?

The commonly held layman’s view of Gaia theory, resulting in part from its emotive name and New Age associations, was that Gaia theory meant that the Earth was alive. This misconception caused problems in the scientific community despite the fact that none of the scientists involved actually believed it. Gaia theory assigned to the Earth system certain life-*like* properties such as global transport and regulation, but the nearest scientific viewpoint to saying that the Earth was alive was the claim that it might be a superorganism (an organism consisting of many smaller organisms each of which is not capable of surviving alone for any significant period of time, such as an ant colony).

Lovelock himself proposed a less contentious way of thinking about the Earth. He viewed the Earth system as an indivisible set of interactions between life and its environment, and had earlier coined a new (less emotive) term to describe this study, ‘geophysiology’, that he hoped would allow for a more scientific approach [120]. The new study of geophysiology was meant to reflect the organism-like properties of the Earth system while pointing out that the Earth system is not an organism in the truest sense.

Further refutations of the ‘Gaia-is-alive’ viewpoint come from two leading Gaia theorists, Lynn Margulis and Tyler Volk. Margulis, one of the founders of Gaia theory and an active proponent of the Gaia hypothesis, argues that Gaia transcends the level of individual organisms:

Gaia, the living Earth, far transcends any any single organism or even any population. One organism’s waste is another’s food. Failing to distinguish anyone’s food from someone else’s waste, the Gaian system recycles matter on the global level. Gaia, the system, emerges from ten million or more connected living species that form an incessantly active body...The sum of planetary life, Gaia, displays a physiology that we recognize as environmental regulation. Gaia itself is not an organism directly selected among many. It is an emergent property of interaction among organisms, the spherical planet on which they reside, and an energy source, the sun...I cannot stress strongly enough that Gaia is not a single organism...[the surface of the planet] behaves as a physiological system in certain limited ways. The aspects that are physiologically controlled include surface temperature and atmospheric composition of reactive gases, including oxygen, and pH. ([126], p.148-154)

Volk also refutes the idea of Gaia as an organism while upholding its organism-like properties:

On the one hand, I experience a delightful sense of being inside a giant metabolism. This perception grows more acute the more I learn, but I am also convinced that Gaia is very different from any organism. Thus I can honestly apply the principles of science to study the global metabolism without postulating a global organism.

What is Gaia? Following Lovelock, I consider Gaia the interacting system of life, soil, atmosphere, and ocean. It is the largest level in the nesting of parts within wholes that encompasses – and thus transcends – living beings, a nesting that ranges from the molecules within cells all the way outward to the Gaian system itself. Like the interiors of organisms, Gaia contains complex cycles and material transformations driven by biological energy. Indeed, Gaia’s inclusion of life means that from some perspectives, it much resembles life. But how Gaia differs from organisms turns out to be its glory.

Consider: Although Gaia has changed through time, it does not evolve in a Darwinian sense. Nevertheless it both contains and is built from evolving organisms. Furthermore, organisms are open, flow-through systems, whereas

Gaia is relatively closed to material transfer across its borders. Gaia exists on its own unique level of operating rules, a level surely as complex as that of organisms and therefore worthy of its own science – which Jim Lovelock calls geophysiology. ([176], p.xiii)

9.2.8 Current status of the Gaia theory

The Second Chapman Conference on the Gaia Hypothesis was held in 2000 with a scope covering:

...not only interactions of biota with atmosphere, the hydrosphere, the soils and the sediments, but also the involvement of biota in maintaining the steady states of key biogeochemical cycles, climate acid/base and redox balances. The three interlinked themes will be Gaia in time, the role of the biota in regulating biogeochemical cycles and climate, and dealing with complexity and feedbacks in the earth system.” ([4], p.6)

The conference confirmed the acceptance of Gaia as a valid field for scientific research; now the focus was on the details of Gaia, how it worked and the mechanisms involved, as opposed to the questions of basic acceptability that had marked the earlier Chapman conference in 1988.

More recently there has been a vigorous debate in the *Climatic Change* journal, where in one special theme issue on Gaia a number of leading Gaian thinkers were asked to give their opinion on the status of the Gaia hypothesis [149]. Kleidon [101] echoes Kirchner in defining a number of competing Gaia hypotheses, but goes further in also providing a metric and a means of testing which of the hypotheses is correct. Kleidon predicted the gross primary productivity of the Earth’s vegetation in different scenarios using simulation models of the interaction between climate and vegetation, and used this data to support his ‘enhancing Gaia’ hypothesis, which stated that life has a beneficial effect on the conditions for life. Lenton [106] also takes a hypothesis-testing approach, considering different explanations for the presence of regulatory feedbacks in the biosphere. He argues first of all that the continuing presence of life on Earth in the face of severe perturbations such as asteroid impacts and volcanic eruptions strongly suggests the existence of regulation, but notes the difficulty in determining whether this regulation is a chance event or statistically probable. Kirchner [98] attacks Gaia theory (on similar grounds to his views noted above), saying that it is far more probable that life has adapted to its environment than the other way around. Finally, Volk [177] gives limited support to Kirchner while remaining supportive of some forms of Gaian mechanism. Interestingly, Volk calls

for new models to be constructed, citing the Guild model of Downing and Zvirinsky [55] as a good example of the direction new modelling work should take.

This was followed closely by a further special issue in which the authors were invited to respond to each others work [150]. Papers were presented by Lenton and Wilkinson [111], by Volk [178], and by Kirchner [99]. We will not cover the contents of these papers here, noting only that there was little fundamental change in the theoretical positions of the contributors. The vigorous continuation of the debate demonstrates that not only had Gaia theory become accepted as a part of mainstream science, but that it remained a controversial topic, with its form and details still the subject of much disagreement and research effort. The general and specific details of the Gaia theory remain to be elucidated.

9.3 Models of Gaia

As mentioned above, there is a genuine need for models in Gaian research [106, 177]. There are significant problems in measuring phenomena on the scale of Gaia, due to the large size and complexity of the Earth system and also to the fact that we, as observers, are a part of the system to be observed. Models allow us to simplify the real world to manageable levels of complexity and to test the validity of assumptions.

This section will look at the various models of Gaia that have been presented in the literature. This will focus mainly on Daisyworld models, as indeed does the Gaian literature, together with brief coverage of Downing & Zvirinsky's Guild model [55].

9.3.1 Daisyworld

The original Daisyworld model [117, 118, 183] was presented in 1983 as a refutation of the teleology criticism of Gaia theory put forward by Doolittle [52] and Dawkins [40]. It did so by demonstrating that simple ecological competition between black and white daisies could regulate the temperature of a fictional planet without the need for any foresight or planning. Daisyworld showed that decentralised control could result in global regulation.

Since its initial appearance, Daisyworld has spawned a large number of papers that extend the basic scheme in order to approach a variety of different questions [3, 35, 41, 42, 93, 105, 107–109, 119, 121, 123, 125, 135, 142, 148, 157, 159, 161, 180, 182, 184, 185, 190, 193]. The Daisyworld model has provided a concrete test case for many aspects of Gaia theory, made popular (despite its simplicity) due to the difficulty of testing Gaia hypotheses in the real Earth system. In the absence of practical real world experiments capable of conclusively testing the validity of the Gaia hypothesis, it seems that Daisyworld has become a proxy for Gaia; if Daisyworld succeeds so does Gaia, and vice versa. While Daisyworld

can help to answer a number of questions about Gaia, it should be remembered that it was never intended to be a complete model of the Earth system and that the Gaia theory does not depend upon it.

Here we will review the original Daisyworld model, before looking at some of the most interesting extensions to it.

9.3.1.1 The original Daisyworld model

The original Daisyworld [117, 118, 183] is a fictional planet inhabited only by two species of daisy. These daisies differ only in their pigmentation: one species is black and the other is white. Solar insolation strikes the planet at a rate that increases slowly over geological time; the warming effect of the insolation is determined by how much is absorbed by the planet and how much is reflected back to space. This in turn is determined by planetary albedo. The albedo of the planet is determined by weighted contributions from the regions of the surface area covered by black daisies, white daisies, and bare earth. Daisy coverage therefore has an effect on the temperature of the planet.

All daisies grow best at a universal optimum temperature and less well as their local temperature moves away from this optimum. The local temperature of a daisy is affected by its albedo (reflectivity) and by the ambient temperature of the whole planet. Black daisies have lower albedo than bare earth and therefore reflect less of the solar insolation. Hence black daisies experience local temperatures that are warmer than bare earth. Conversely, white daisies have higher albedo than bare earth and reflect more of the sun's light, so white daisies experience local temperatures cooler than bare earth. The effects of warming by black daisies and cooling by white daisies allow global temperature to be altered in both directions by the biota.

The formulation for the basic Daisyworld is presented below, taken from Watson and Lovelock [183]. Note that some aspects of the model are perhaps over-complicated for the phenomenon that Watson and Lovelock were trying to demonstrate. For example, there is no need for such a realistic implementation of heat radiation in Equation 9.4, where what is important is some heat loss gradient from the planet to outer space. Similarly there is no need for such precise values for the constants in Equation 9.3, where all that is required for the model to work is any form of parabolic function. These details mask the true simplicity of the Daisyworld model and seem irrelevant given the arbitrary nature of other features of the model (such as genetic variation). We can only speculate at the reasons for their inclusion when the aim of the model was to show how global regulation could emerge without teleology or central control, but this should not be allowed to devalue the valuable contribution the Daisyworld model has made to the Gaia debate.

Equation 9.1 describes the change over time in the daisy populations, expressed as the proportion of the surface of the planet they cover. A daisy population expands to cover more of the surface area of the planet at a rate directly proportional to its current size and to the amount of bare earth available, and decreases at a constant rate due to daisy death.

$$\begin{aligned}\frac{d\alpha_w}{dt} &= \alpha_w(x\beta_w - \gamma) \\ \frac{d\alpha_b}{dt} &= \alpha_b(x\beta_b - \gamma)\end{aligned}\quad (9.1)$$

where α_w and α_b are the proportion of the fertile surface area covered in white and black daisies respectively. x is the proportion of the fertile surface left bare. β_w and β_b are the growth rates of white and black daisies per unit time. γ is the death rate per unit time (constant for both types of daisy).

Equation 9.2 states that all of the fertile surface area of the planet is either uncolonised, or colonised by black or white daisies. The total fertile surface area of the planet is taken as a dimensionless constant P , set to unity. Then the proportion of the total fertile surface area that is uncolonised is (x), found as the total area P minus the proportions covered by black (α_b) and white (α_w) daisies.

$$x = P - \alpha_w - \alpha_b \quad (9.2)$$

Equation 9.3 defines the growth rate of the daisies as a parabolic function of their local temperature (T_i , where i denotes the colour of the daisy), centred on an optimum rate of 1 at $T_i = 22.5^\circ C$ and falling to zero by $T_i = 5^\circ C$ and $T_i = 40^\circ C$.

$$\beta_i = \begin{cases} 1 - 0.003265(22.5 - T_i)^2 & : \quad 5 < T_i < 40 \\ 0 & : \quad otherwise \end{cases} \quad (9.3)$$

Equation 9.4 states that the radiation emitted by the planet must equal that absorbed. ρ is Stefan's constant, T_e is the effective planetary temperature ($^\circ C$), L is the (dimensionless) solar luminosity, and S is a constant (with units of flux).

$$\rho(T_e + 273)^4 = SL(1 - A) \quad (9.4)$$

Equation 9.5 gives the mean albedo of the planet (A) as the sum of the albedos (A_i) of bare earth, black, and white daisies, weighted by the proportions (α_i) of the surface area they cover (where $i = g$ denotes bare earth and $i = b$ and $i = w$ denote black and white daisies respectively). For the simulations performed $A_g = 0.5$, $A_b = 0.25$ and $A_w = 0.75$, so that $A_b < A_g < A_w$.

$$A = A_g \alpha_g + A_w \alpha_w + A_b \alpha_b \quad (9.5)$$

Equation 9.6 deals with heat flow between patches in a simple fashion by relating local temperatures (T_i) to the effective planet temperature (T_e) by albedo and a constant (q) that measures the amount of redistribution of solar energy. If $q = 0$ then daisy patches have the same temperature as the planet, while if $q > \frac{SL}{\rho}$ heat flow is (impossibly) against the temperature gradient. Watson and Lovelock set $q = 2.06 \times 10^9$, well below the critical value. Equation 9.6 preserves the energy balance in the Daisyworld system.

$$(T_i + 273)^4 = q(A - A_i) + (T_e + 273)^4 \quad (9.6)$$

The system of equations describing Daisyworld were demonstrated by Watson and Lovelock to converge to a single steady state attractor for any given parameterisation, no matter what the initial condition. Watson and Lovelock used a computer to numerically integrate the Daisyworld system to a steady state attractor to give equilibrium values for T_e for different values of L . L was incrementally increased to represent a gradual increase in solar luminosity similar to that experienced by the Earth as the Sun has grown hotter during Earth's history. Both black and white daisy seeds were assumed to be present in the soil of Daisyworld, so that either species could spontaneously germinate if conditions were suitable.

Figure 9.1 plots the steady state response of the Daisyworld system as L is incremented. The constants in the system were set as follows: $P = 1$, $\gamma = 0.3$, $S = 9.17 \times 10^5$ for convenience (since this makes L approximately equal to 1), $A_g = 0.5$, $A_b = 0.25$, $A_w = 0.75$. Figure 9.1 shows that daisies with neutral albedo have no effect on the effective temperature of the planet, despite maintaining a large population when conditions are suitable (Figure 9.1(a)). Black daisies have a warming effect (Figure 9.1(b)), while white daisies have a cooling effect (Figure 9.1(c)). Black and white daisies together maintain temperature around the optimal level for growth by a process of ecological competition (Figure 9.1(d)).

Initially no daisies are present, then when increasing luminosity warms the planet to a sufficient level black daisies are able to grow because they have a local warming effect. The black daisy population quickly spreads and warms the planet towards the optimal level for growth, allowing white daisies to grow also. The black daisy population initially dominates because the local warming effect of black daisies allows them to out-compete the white daisies in cold conditions. As luminosity increases, the black daisy population starts to get out-competed by the white daisies, whose local cooling effects become more beneficial, and eventually the white daisy population dominates. The black

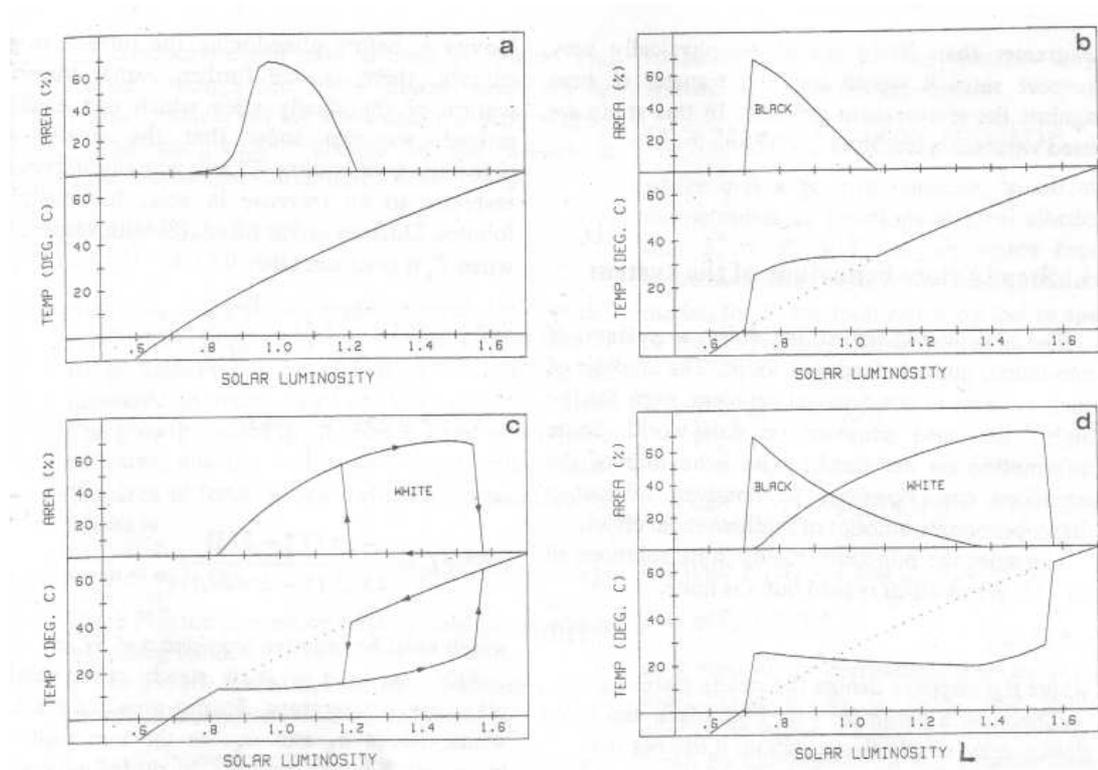


Figure 9.1: Results from original Daisyworld. Percentage of fertile area covered by black and white daisies and effective temperature are plotted against increasing solar luminosity. Solid and dotted lines show effective temperature of planet with and without life. [a] Neutral daisies (albedo 0.5, equivalent to bare ground) have no effect on temperature, dotted and solid lines are coincident. [b] Black daisies (albedo 0.25) only can regulate temperature upwards towards optimal level. [c] White daisies (albedo 0.75) only can regulate temperature downwards towards optimal level. Also shown is effect of decreasing solar luminosity. [d] Competition between black and white daisies leads to temperature regulation. All results reproduced with permission from Watson and Lovelock [183]; poor quality of reproduction is the fault of the current author.

daisy population goes extinct when luminosity warms the planet to a point where the additional warming effect of the black daisies makes them too hot to survive. Eventually the white population also goes extinct when its cooling properties are insufficient to keep its local temperature viable for growth. At this point planetary temperature regulation collapses and there is a sudden large increase in temperature as the planet returns to its abiotic stable state.

9.3.1.2 Ecology in Daisyworld

Lovelock [121] experimented with the ecology of Daisyworld by looking at the case where there are 100 daisy species, each coloured a different shade of grey to give a range of albedos from 0.25 (black) to 0.75 (white). He found that temperature regulation was unaffected, but that competitive exclusion meant that only one or two daisy species co-existed at any particular level of solar luminosity, a result that had earlier been shown by Maddock [125]. Von Bloh et al [180] also showed that competitive exclusion reduced the number of co-existing species in their 2-dimensional cellular automata Daisyworld, but that this actually improved temperature regulation.

Lovelock extended this model to include further trophic levels in the form of herbivores (whose growth depended on the size of the daisy population) and carnivores (who preyed on the herbivores). He showed that biodiversity in Daisyworld was greatest just after perturbation and was least after a long period without perturbation when the system approached a steady state.

Harding and Lovelock [76] also added herbivores, but allowed a variety of different herbivore feeding strategies. They found that temperature regulation was unaffected, but that the types of herbivore feeding strategy in the system influenced the trajectory of temperature regulation displayed by the system. A later paper showed that when multiple herbivore strategies were present, the dominant strategy was determined by the dynamics of daisy-environment interaction [75].

Cohen and Rich [35] presented results showing that in some circumstances the presence of life on Daisyworld had a destabilising (rather than stabilising) effect on global temperature. They implemented stronger competition between daisy species by reducing the growth rate of a species dependent on the size of the population of the other species, and showed that for some changes in luminosity the system displayed a larger change in temperature with life present than with life absent.

On the whole increasing the complexity of the ecology in Daisyworld seems to make little difference to the fundamental behaviour of the model. Adding more daisy species does not effect regulation, while adding further trophic levels does not seem to lead to

any new insight. The most interesting findings in this section are the occurrence of competitive exclusion in the Lovelock's multi-species model [121], and Cohen and Rich's demonstration that when competition between species is made extreme, life can actually have a destabilising effect on global temperature [35].

9.3.1.3 Chaos in Daisyworld

Zeng et al [193] apparently showed that the surface temperature in Daisyworld could display chaotic behaviour. They used discretised versions of the the equations given above, together with a forward difference numerical integration algorithm with a timestep $\Delta t = 1$. Zeng et al claimed that the chaotic behaviour that resulted not only invalidated the results achieved by Watson and Lovelock [183], but cast doubt on the Gaia theory as a whole.

However, Jascourt and Raymond [93] pointed out that the discrete and differential versions of the Daisyworld are fundamentally different systems. Jascourt and Raymond show that the large timestep used by Zeng et al corresponds to a one-generation lag in the population and environmental update equations, which is the cause of the chaos observed by Zeng et al. Furthermore, Jascourt and Raymond show that some of the parameterisations used by Zeng et al to show chaotic behaviour actually result in impossible negative population sizes at various points in the chaotic trajectory. Finally, Jascourt and Raymond point out that even in the chaotic regime of the discrete system, the mean temperature in the presence of daisies is stable and that the results of Lovelock and Watson [183] are fundamentally supported.

A further explication of the confusion over chaos in Daisyworld is provided by De Gregorio et al [41,42], who introduce a time lag into the differential Daisyworld model of Lovelock and Watson [183]. De Gregorio et al demonstrate that as the size of the time lag is increased, the system moves from fixed point to limit cycle and eventually to chaotic solutions, supporting the findings of Jascourt and Raymond [93].

A footnote to the debate over chaos in Daisyworld is supplied by Lenton and van Oijen [109], who point out that there is no biological justification for the delays or time lags that give rise to periodic or chaotic solutions.

9.3.1.4 Spatial Daisyworlds

The original Daisyworld model is dimensionless. Although there are variables representing the areas covered by black and white daisies, and the remaining bare earth, there are no spatial aspects to the way the model works other than an assumption that black and white daisies live in patches large enough to have their own local temperature.

Von Bloh et al [180] presented a cellular automata model of Daisyworld that modelled

the surface of the planet as an array of patches, each of which could be populated by a daisy species or left bare. They incorporated lateral heat flow between adjacent patches and showed that this allowed the coexistence of multiple daisy species (as compared with the competitive exclusion seen in non-spatial Daisyworlds), since patches containing mal-adaptive daisy species with albedo too high or too low for the current level of insolation could be warmed or cooled by neighbouring patches populated by better adapted daisy species. Von Bloh et al showed that global temperature regulation in the face of increasing solar luminosity was more effective in their 2-dimensional CA than in the non-spatial equivalent.

Other spatial Daisyworld models are reported on elsewhere in the literature [3, 109], where one interesting result is that spatial models allow ‘desert’ formation at certain levels of solar luminosity [3].

9.3.1.5 Evolution in Daisyworld

The original Daisyworld model includes evolution only in the limited sense that the daisy population most suited to the environment will grow most quickly, but there have been many extensions to the model to try and address the question of whether the emergent regulation displayed by Daisyworld (and by extension Gaia in general) is compatible with evolutionary theory.

One criticism of Daisyworld (and Gaia) by evolutionary theorists is that the global regulation will be vulnerable to cheaters, which benefit from regulation but do not contribute. To test this hypothesis, Lovelock developed a variation of the original Daisyworld model that added a cost for producing pigmentation expressed as a growth rate reduction [119]. This model included grey daisies that did not produce black or white colouring and thus did not incur this cost. Lovelock showed that the presence of these potential ‘cheaters’ did not destroy regulation, since the grey daisies only had a competitive advantage at certain temperatures. This result refuted the cheater criticism of Daisyworld and was later reconfirmed [108].

Lovelock’s multiple-species model extended the size of the ‘gene pool’ by including 100 daisy species with different albedos and showed that regulation still occurred in this scenario [121]. Stöcker [159] then showed that when a more realistic mechanism for mutation of daisy albedo was used regulation not only still occurred, but was more efficient. Lenton [105] noted the criticism of Keeling [96] concerning the fact that in these models all possible daisy species are pre-specified, and showed that regulation still occurred in a model where albedo mutated by occasional small perturbations to value passed down from parent to offspring which were then faithfully replicated. Other papers have also

confirmed that mutation of albedo does not destroy regulation and in some cases may improve it [108, 109, 180].

The other aspect of daisy evolution that has been addressed is adaptation of the optimum growth temperature. Saunders [148] showed that when the optimum temperature of black daisies was switched to be 22.5°C and the optimum temperature for white daisies switched to 17.5°C (as might be expected to occur if the daisies' metabolism adapted to prefer the local temperature created by their respective albedos), planetary temperature regulation still occurred, albeit with a smaller range of stability. This finding was contradicted by Robertson and Robinson [142], who used spreadsheet simulations to show that temperature regulation did not occur if the preferred growth temperature of the daisies was allowed to mutate freely. Lenton and Lovelock [107] responded to this finding by pointing out the existence of constraints on evolutionary adaptation imposed by physical and chemical laws, and then demonstrating that regulation occurred when constraints on the mutation of preferred growth temperature were enforced.

An interesting result was shown by Wood et al [190], who found that when mutation of both growth temperature and albedo was incorporated into a spatial Daisyworld model with heat flow, one possible result was the occurrence of stable oscillations with a period of hundreds of daisy generations. The evolution of growth temperature interacts with the evolution of albedo, so that low temperature/high albedo species could co-exist with high temperature/low albedo variants. Global temperature regulation still occurred, but as an average of a stable oscillation in the habitable range.

9.3.1.6 Cut-down Daisyworld

Harvey [80] presented a 'cut-down Daisyworld' model in which he simplified the mathematics of the system as far as possible while retaining its regulatory properties. Harvey replaced the parabolic growth function of the daisies in the original Daisyworld model by a piecewise linear 'hat' function (see Figure 10.1 in Chapter 10) and also replaced the semi-realistic heat radiation and transfer functions of the original model with simplified linear caricatures of the underlying physics. The interaction between different daisy species was restricted to indirect interaction by heat transfer between daisy patches; no competition for space or mutual inhibition was included.

Harvey found that global temperature regulation was easily displayed in the 'cut-down Daisyworld' system, and that its occurrence depended on the existence of the 'hat' function (although the precise shape of the hat was irrelevant and could be any peaked shape) and on an intermediate level of heat flow between the black and white daisy patches. If the two patches were too tightly coupled, they acted as a single mixed-species patch in

which neither species was able to set up an advantageous local temperature, and the global regulation broke down. If they were not coupled at all, it was as if there were two separate planets and the concept of collective regulation had no validity.

9.3.2 Other Gaia models

Daisyworld models have been the central modelling paradigm for studies of Gaia theory, but additional models from an artificial life perspective have been developed by Downing [53–55]. The Euzone model [53] was a model of the evolution of aquatic ecosystems and their effect on the chemical environment, but no regulation was demonstrated. The Guild model [55] was an attempt to reconcile Gaia theory and natural selection, in which the evolution of nutrient recycling guilds and environmental control was explored. The Metamic model [54] incorporated a more realistic chemistry than Guild and was developed in order to test the logical plausibility of Gaian phenomena occurring in a system following basic rules governing chemical reactions, metabolism, and energy transfer. The Guild model is briefly described below.

9.3.2.1 The Guild model

The Guild model [55] looks at an evolving ecosystem of many species of bacteria-like organisms that interact via a shared chemical environment. Guild considers not only environmental regulation by the biota but also Volk's proposed Gaian metric of nutrient cycling [176]. Guild was intended to support and extend Daisyworld by allowing a much greater range of genetic variation and setting an environmental control problem that could not be answered by a single species acting alone; in Guild the chemical environment is regulated as a result of the collective metabolic activity of the biota.

Guild is an individual-based evolutionary simulation model. Organisms in Guild consume and excrete chemical nutrients in genetically specified ratios. An organism cannot consume what it excretes. At each time step each organism will attempt to consume an amount of nutrients A_{feed} , depending on availability, where A_{feed} is given by Equation 9.7.

$$A_{feed} = (X)^{0.75} r_f S \quad (9.7)$$

Here X is the current biomass of the organism, r_f is a universal constant base feeding rate, and S is the satisfaction of the organism with its local environment. S is an error function of the relative proportions of the organism's local chemical environment compared to a pre-determined optimum ratio that is universal to all organisms. S is given by

Equation 9.8.

$$S = e^{-k_{sat}\epsilon_{sat}} \quad (9.8)$$

where k_{sat} is a universal satisfaction coefficient that determines the strength of influence of satisfaction on the amount of food consumed, and ϵ_{sat} is a measure of the deviation of the organism's local environment from the pre-defined optimum given by Equation 9.9.

$$\epsilon_{sat} = \frac{1}{n} \sum_{i=1}^n \left| \frac{E_i^o - E_i'}{E_i^o} \right| \quad (9.9)$$

where E_i^o is the optimum proportion in the environment of the i^{th} chemical nutrient and E_i' is the effective proportion of this chemical experienced by the organism. These quantities are calculated from the consumption and excretion ratios of the organism and from the current chemical constitution of the global environment by Equation 9.10.

$$E_i' = E_i(1 + k_{filter}(F_{i,out} - F_{i,in})) \quad (9.10)$$

where E_i is the level of the i^{th} chemical in the global environment, $F_{i,in}$ and $F_{i,out}$ are the proportions of this chemical that the organism consumes or excretes respectively, and k_{filter} is a universal constant determining the degree to which organism's can filter their perception of the chemical environment.

There are a number of other aspects to the Guild model concerning reproduction, mutation, death and recycling of dead material that are not covered here. Here we have only given the key equations that are needed to understand how the model operates.

Organisms gain biomass at a rate dependent on availability of nutrients and their current satisfaction with their local environment. Satisfaction varies inversely with deviation of the organism's local environment from a pre-defined optimum. The local environment of the organism is determined by the global environment 'filtered' by the effect of the organism's consumption and excretion of chemicals. Thus an organism can create a local buffer against the global environment, which may offer selective advantage over other organisms by promoting growth.

Figure 9.2 shows the result of a typical simulation run. The model is seeded with a single species and is parameterised so that there are 4 chemical nutrients $\{N_1, \dots, N_4\}$ with an optimal growth ratio of $\langle E_1^o, E_2^o, E_3^o, E_4^o \rangle = \langle 0.4, 0.3, 0.2, 0.1 \rangle$. Population size, nutrient cycling ratio (roughly speaking the amount of the nutrient consumed divided by the influx of that nutrient per timestep), and environmental nutrient fractions (the proportions of the different chemical nutrients present in the environment), are plotted against generational time.

Initially the single-species population can only exploit a single nutrient and biotic growth is constrained by the amount of that nutrient entering the environment at each timestep. Eventually new species are created by mutation that can use other nutrients and there is a population boom. Nutrient cycling networks are formed and the biota start to regulate the chemical environment.

The Guild model shows that the key Gaian metrics of regulation and nutrient recycling can evolve in a simulated microcosm. There is no central control, and only standard individual-level selection is used, yet coordinated distributed control evolves. Thus the results of the Guild model support the results from Daisyworld and show that regulation can still emerge when the genetic space is large.

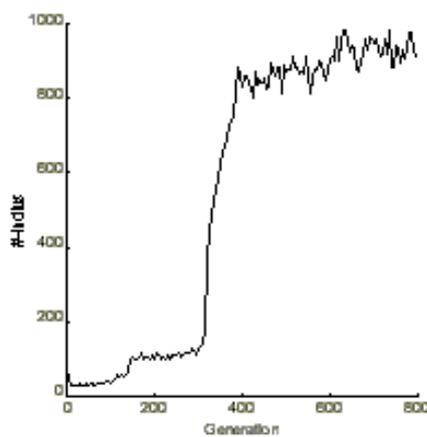
9.4 Aims, objectives, questions

We have seen that the modern Gaia theory posits that the interaction between life and its environment is bi-directional, and that life can to some extent regulate its environment to conditions suitable for life. The central debate over Gaia theory concerns the compatibility of Gaia theory with evolutionary theory. Can regulatory feedbacks evolve? Are they vulnerable to cheaters?

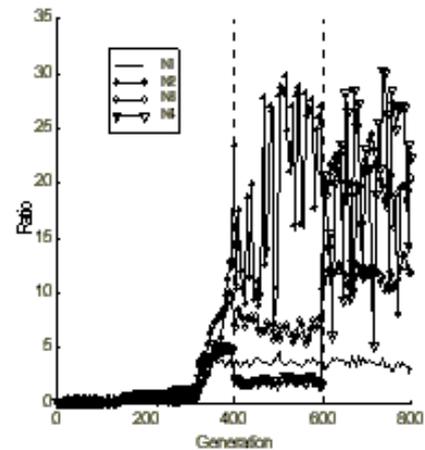
Daisyworld, supported by Guild, has answered a number of questions in the Gaian debate. Daisyworld has been used to refute the teleology criticism that was previously aimed at the Gaia hypothesis, and has also provided limited evidence that Gaian regulation can evolve. Guild has shown that regulation can emerge as a result of the ecological and evolutionary interactions between simple organisms in a simulated microcosm, thus showing that regulation can occur as a result of ‘by-products’ rather than as a costly process additional to metabolism.

However, certain kinds of evolution destroy regulation in Daisyworld; when the preferred growth temperature can mutate freely regulation collapses. While there have been several efforts to incorporate evolution into Daisyworld, there are still questions to be answered, such as what happens when both albedo and growth temperature evolve at the same time, and what effect different mutation rates would have. The original Daisyworld model incorporating evolution is not easy to analyse, and it seems that there would be utility in developing a simple and analytically tractable model that allows the testing of different hypotheses concerning the effect of evolution on temperature regulation.

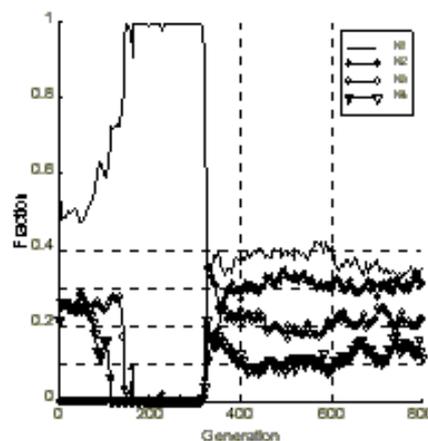
Also, it is not clear that Guild is an entirely distinct model from Daisyworld. While Guild uses different language and models distributed environmental control in an ecosystem, both Daisyworld and Guild share some key assumptions. One is that individual organisms can create a local buffer against the larger environment that affords them some



(a) Population is initially small but expands rapidly once mutation allows the ecology to take advantage of previously unused nutrients.



(b) Nutrient cycling is poor until population growth allows sufficient diversity for cycling loops to be formed.



(c) Environmental chemical ratios are not regulated until population growth and diversity allows formation of a stable food web that regulates environmental chemicals close to the optimal ratio.

Figure 9.2: Results from a Guild simulation run. Vertical axes show: (a) Population size, (b) nutrient cycling ratio, (c) environmental chemical ratios. Horizontal axes in all plots show number of generations elapsed. Population is initially small, until mutation to a metabolism that can capitalise on stored environmental nutrients allows a population boom. An ecology develops and nutrient cycling ratios in the ecosystem rise dramatically. At the same time the biota establish control of the levels of different chemicals in the environment; the onset of regulation. Figures reproduced with permission from [55]; poor quality of reproduction is the fault of the current author.

selective advantage. Another is that the only evolutionary adaptations that are beneficial at an individual level are also those that contribute towards Gaian regulation; there is no possibility for selfish mutations that have an adverse effect on regulation. It would be interesting to look at the effect of relaxing these assumptions.

With these observations in mind, we can now frame some questions to be answered in subsequent chapters.

1. Under what conditions will regulation occur in an evolutionary Daisyworld?
2. Does regulation occur when organisms cannot create a local buffer against the environment?
3. Does regulation occur when selfish mutations do not necessarily contribute to it?

To explore the answers to these questions, we will first of all develop a simple Daisyworld model that allows for straightforward incorporation of different kinds of evolutionary adaptation. We will use this model to find the conditions under which regulation of global temperature will evolve, and to clarify our understanding of the compatibility of evolution with Gaian regulation in Daisyworld. Next we will develop a new model that seeks to relax the common assumptions of the Daisyworld/Guild models and advance the theory of environmental regulation by the biota, hoping to gain new insight into the central question of Gaia theory: can Gaian regulation evolve in a manner consistent with evolutionary theory?

Chapter 10

Evolution in Daisyworld

10.1 Overview

As set out at the end of the preceding chapter, the aim of this chapter is to look at the conditions under which environmental regulation will occur in an evolutionary Daisyworld model. We approach this question by constructing a simplified Daisyworld model that allows easy implementation of daisy adaptation while maintaining similar dynamics to the original model. We take inspiration from Harvey's 'cut-down Daisyworld' [80], but further simplify the model and extend it into a 2-dimensional cellular automata model that is more amenable to the inclusion of evolutionary adaptation. The simplicity of the model aids understanding of the mechanisms underlying the evolution of regulation in Daisyworld models. We look at adaptation of daisy albedo and of the preferred growth temperature, starting by replicating known results from the literature and then moving on to present some new results concerning the importance of constraints for the evolution of regulation.

The chapter begins with a description of the model used, followed by presentation of results.

10.2 Rationale

The original Daisyworld model [183] incorporates two species of daisy, identical except that one is black (with low albedo) and the other is white (with high albedo). Daisy

albedo alters the local temperature of each daisy patch, with daisies assumed to live in single-species clumps large enough to maintain their own local temperature. This in turn alters the growth rate of the daisies, which varies as a function of temperature. Since the albedo of bare earth lies between the albedos of black and white daisies, population dynamics allow global temperature to move away from that expected of a dead planet. Competition between black and white daisies led to global temperature regulation around the optimal temperature for daisy growth; deviations away from this point were counteracted by negative feedback engendered by the selective advantage gained by one of the daisy species away from this point. Black daisies out-compete white daisies at low temperatures because of their ability to increase local temperature, and vice versa at high temperatures. Regulation was observed for a significant range of solar luminosity, outside which the planet was too cold or too hot to support daisies of any colour.

The Daisyworld model used reasonably accurate approximations of the real-world phenomena on which it was based and thus incorporated quite complicated mathematical formulations of (for example) the interaction between solar luminosity and the level of heat radiation emitted by the planet. Harvey [80] developed a simplified model, his ‘cut-down Daisyworld’, that used much simpler approximations but conserved the essential regulatory behaviour of the system. Following Harvey and simplifying even further, we present the very basic model described below.

10.3 Model

Our model is a cellular automata model in which patches are arranged in a 2-dimensional toroidal lattice (another CA Daisyworld model was presented by von Bloh et al [180], but in a different form and with different aims). Each patch may be barren (bare earth) or may contain a single species of daisy. Barren patches can be colonised by daisies from neighbouring patches, while living patches may die. Each patch has a local temperature that changes in relation to solar luminosity (applied at an equal level to all patches) and to its albedo (determined by the presence of daisies). The global temperature of the planet is taken as the mean of all the local patch temperatures. This scheme is covered in more detail below.

10.3.1 Daisies

A daisy species is represented by an albedo and a growth function, i.e., each daisy has a colour and a preferred growth temperature. The daisy population can be thought of as points scattered in a 2-dimensional gene space. Albedo is drawn from the range

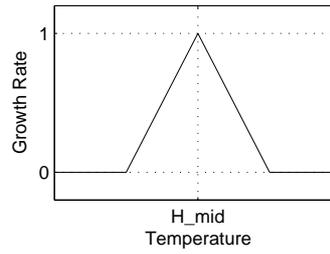


Figure 10.1: An example growth function. The growth rate of all daisy species varies from 0 to 1 as a piecewise linear function of temperature.

[0.25, 0.75] (representing a continuum from black to white), while the growth function is a piecewise linear function of local temperature that has the qualitative form shown in Figure 10.1 and is given by Equation 10.1. This ‘hat-shaped’ function has the same width 2δ at its base for all daisies, so the hat function for a single daisy species can be represented by the location of its centre point H_{mid} .

$$G = \begin{cases} 0 & : & T < (H_{mid} - \delta) \\ \frac{T - (H_{mid} - \delta)}{\delta} & : & (H_{mid} - \delta) \leq T < H_{mid} \\ \frac{(H_{mid} + \delta) - T}{\delta} & : & H_{mid} \leq T < (H_{mid} + \delta) \\ 0 & : & T \geq (H_{mid} + \delta) \end{cases} \quad (10.1)$$

where G is the growth rate of the daisy species, T is the current local temperature of the patch, H_{mid} is the centre point of its growth function and 2δ is the width of the growth function at its base. In all the experiments reported here $\delta = 15$ so that the hat function reached zero at $H_{mid} \pm 15$.

Daisies are assumed to either fully occupy a patch or not to be present. The growth rate of a daisy species determines its likelihood of colonising a neighbouring bare patch. High growth rates lead to increased colonisation and the spread of the species.

10.3.2 Seeding

An empty patch may be seeded with a new daisy species with low probability (0.03 in the simulations described below). When seeding occurs an entirely new daisy type is randomly generated from the set of permissible values for albedo and growth function parameters. Seeding allows new genetic stock to enter the world and takes the place of mutation in the evolutionary process.

10.3.3 Colonisation

Empty patches may be colonised by daisy species living in neighbouring patches. Each neighbour species has a chance to colonise that is proportionate to its growth rate. This is implemented by assigning a probability $P(C_i)$ to the event C_i that the empty patch is colonised by the i^{th} neighbouring patch (alive or dead). This is expressed by Equation 10.2 below.

$$P(C_i) = \frac{G_i}{N} \quad (10.2)$$

where G_i is the growth rate of the i^{th} neighbouring patch and N is the total number of neighbours. Note that the growth rate of a dead patch is zero. Thus daisies with a higher growth rate have a higher likelihood of colonisation. Also, a daisy species occupying multiple neighbouring patches has a higher likelihood of colonisation due to having more ‘tickets in the lottery’.

10.3.4 Death

If a daisy species living in a patch has a growth rate of zero, it is assumed not to be able to survive and the patch becomes empty. Also, daisies living in a patch will die (and the patch become empty) with a probability of 0.1 at each timestep. This may be seen as a simple instantiation of death by natural causes and serves to promote selection and competition.

10.3.5 Calculation of patch temperature

Local patch temperature depends on the current temperature of the Sun (traditionally taken in Daisyworld models as a monotonically increasing value), the albedo of the patch (determined by daisy growth), and heat loss to space. The rate of change of local patch temperature is therefore given by Equation 10.3 below, where T_P is the patch temperature, T_S is the temperature of the Sun, and α is the patch albedo.

$$\frac{dT_P}{dt} = (1 - \alpha)(T_S - T_P) - T_P \quad (10.3)$$

In Equation 10.3, the first term increases patch temperature in proportion to the heat gradient ($T_S - T_P$) from the Sun to the patch, moderated inversely by patch albedo α so that high albedo reduces heat flow (by reflecting solar insolation back to space). The second term reduces patch temperature to account for heat loss to space by radiation and

is proportional to the negative heat gradient between space and the patch ($0 - T_p$), where the temperature of space is assumed to be zero.

Equation 10.3 does not include a term for heat flow between patches. It was found that when heat transfer between neighbouring patches was included in the model, it made no qualitative difference to the results achieved. For this reason it is omitted here for clarity. The exception to this rule is when heat transfer is instantaneous, i.e., when patches equalise temperature instantly. This special case will be considered later in this chapter.

Patch temperatures are integrated numerically using Euler's forward method. The global temperature of the planet is taken as the mean of all the patch temperatures.

10.3.6 Cellular automata update

The results presented below were gathered from a 10×10 toroidal CA where each patch has 4 neighbours at top, bottom, left and right. Runs were also performed with larger (20×20 , 50×50) CAs, and with CAs where each patch had 8 neighbours (added at top-left, top-right, bottom-left, bottom-right), with no qualitative changes in results. The CA is synchronously updated at each timestep (i.e., all patches are updated simultaneously) by testing for colonisation, seeding and death in that order. Although in discrete systems synchronous update can cause artefacts, here states are continuous so this should not be a problem [49,81]. Runs with asynchronous update produced similar results. T_S is typically increased from 100 to 500 in increments of 2, and the CA is updated for 1000 timesteps for each increment in T_S to allow the daisy population to stabilise for the new level of external forcing.

10.4 Results 1: Replication of existing Daisyworld results

First of all we compared the results generated from our model with known results generated from existing Daisyworld models. In all of the following experiments the albedo of bare earth was set to 0.5 and the world was initialised with all patches bare.

The primary Daisyworld phenomenon, that of temperature regulation by competition between black and white daisies [183], was considered first of all. We set the albedo of black daisies to 0.25 and the albedo of white daisies to 0.75. Results are shown in Figure 10.2, which displays global temperature regulation occurring by competition between daisy species as it does in the original Daisyworld model.

The next significant result to be repeated is that allowing albedo to mutate does not affect regulation, and may in some cases actually increase its range [105, 121, 159, 180]. For this scenario we allowed albedo to take any value in the range $[0.25, 0.75]$, corresponding

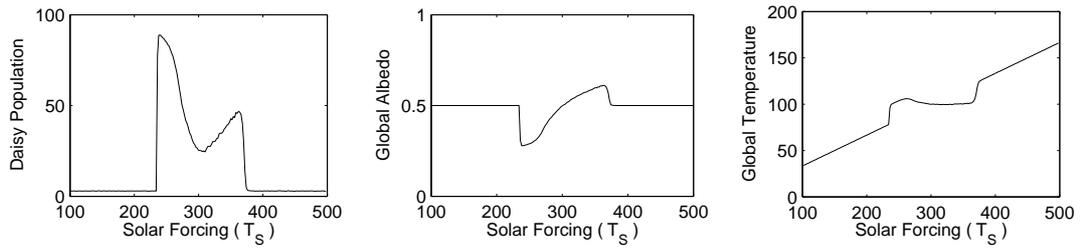


Figure 10.2: Daisy population, global albedo and global temperature for a world with both black (albedo = 0.25) and white (albedo = 0.75) daisies. Temperature regulation occurs as global albedo is adjusted by competition between black and white daisy species.

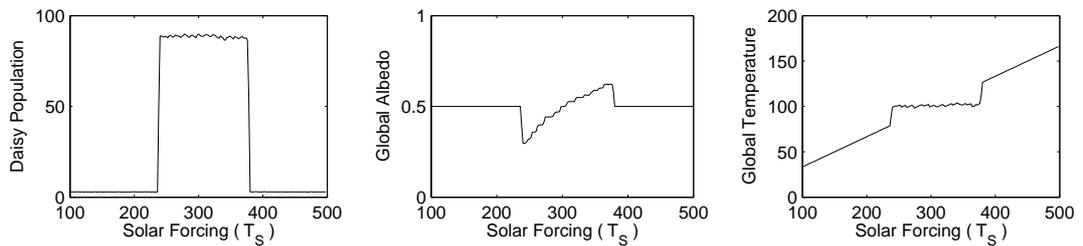


Figure 10.3: Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$. Temperature regulation occurs as albedo of dominant daisy species varies in response to changing solar luminosity.

to the full range from black to white. Temperature was regulated as before, although in this case it is by a steady shift in the albedo of the dominant daisy species to maintain the global temperature close to the optimal level, rather than competition between black and white daisies. The overall effect is the same at a global level; temperature regulation in this case and in the previous case is achieved by keeping the mean global albedo close to the level which keeps temperature optimal. This in turn is a result of selection for the daisy species with the highest growth rates.

Having shown that temperature regulation is not affected by mutation of albedo, the next result is to show that unconstrained mutation of the growth function causes the regulation to break down [142]. Here we do this by allowing H_{mid} to vary freely in the range $[70, 130]$. No regulation is observed, although the daisy population flourishes. The mutation in growth function simply tracks the solar forcing; the daisies adapt themselves to the environment rather than adapt the environment to themselves. The tracking is not precise, as the mutating albedo allows brief periods of quasi-regulation when the population becomes fixated on a particular growth function and uses the albedo to maintain the global temperature this value of H_{mid} requires. In this scenario albedo can be seen as a free variable, since a suitable growth function can be found to give optimal growth for any albedo

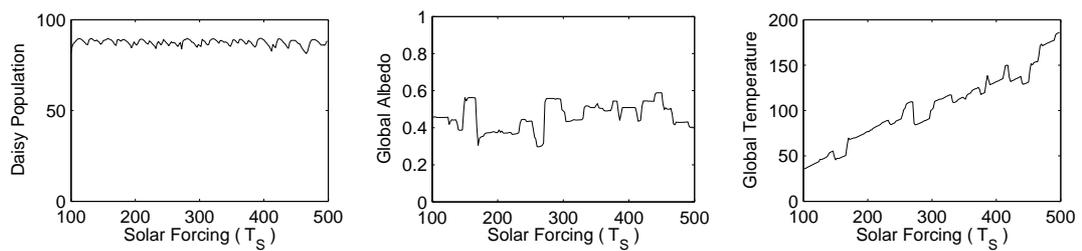


Figure 10.4: Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$ and where the centre of the growth function is allowed to mutate freely. No temperature regulation occurs because daisies can adapt their preferred growth temperature to suit current environmental conditions.

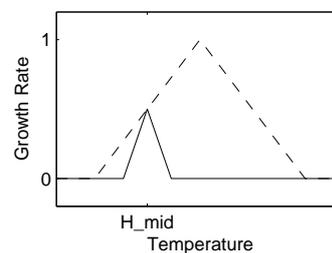


Figure 10.5: Constrained growth function. The maximum achievable growth rate of all daisy species varies from 0 to 1 as a piecewise linear function of temperature. The height of the peak of a species' individual growth function is constrained.

level.

Lenton and Lovelock [107] showed that when there is some constraint on the mutation of the growth function, regulation will again emerge. They set up a Daisyworld model where the centre of the hat function was mutated towards the current ambient temperature, but where the maximum growth rate (i.e., the maximum height of the growth function) fell away to zero with distance from some optimal value, in a way supposed to be analogous with the decline in maximum achievable photosynthesis rate varies with temperature in plants. We implemented this by letting the maximum growth rate decline linearly to zero with distance from an optimal temperature of 100 (see Figure 10.5). We observed similar results to Lenton and Lovelock [107], in that regulation was observed to occur, but with a more gradual tailing in and tailing out than with the non-evolvable growth function.

We have now shown similar results to the most significant results achieved with more conventional Daisyworld models. Regulation of global temperature has been shown to occur when there is some constraint on evolution of the growth function (itself a very well-supported assumption based on empirical evidence from biology given that life has only been found within certain bounds of temperature, acidity, etc.), and where the daisies have some means by which they can influence their local environmental temperature. If

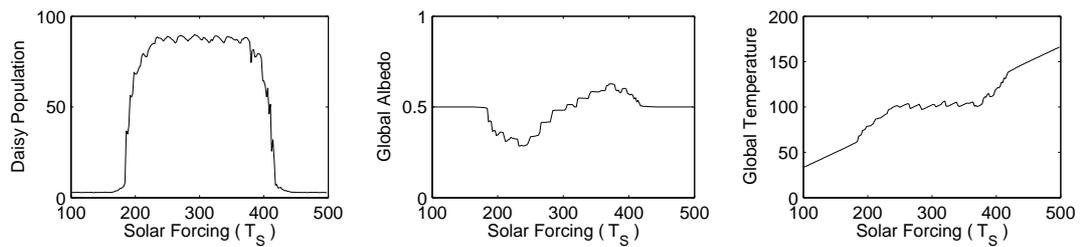


Figure 10.6: Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$ and where the centre of the growth function is allowed to mutate freely. Maximum achievable growth rate declines linearly with distance from $T = 100$, reaching zero at $T = 100 \pm 30$. Temperature regulation occurs because preferred growth temperature is constrained and daisies cannot always adapt their preferred growth temperature to current conditions, allowing selective advantage to be gained from using albedo to alter local temperature.

there are no constraints on the growth function (i.e., no *need* to regulate) or global albedo cannot vary (i.e., no *means* to regulate) then regulation breaks down.

For completeness, we have also run the model with heat transfer between neighbouring patches [180], and found that the qualitative nature of the results is unchanged for all of the above scenarios.

10.5 Results 2: Constraints on evolution and their implications for environmental regulation

It seems that the key criteria for regulation of global temperature to emerge are *need* and *ability*. Unless there is some reason for the daisies to alter their local environment, i.e., some selective advantage to be gained from doing so, then regulation will not occur. If evolution is added to the model, then the only cases in which daisies have a reason to alter their environment are those in which the evolutionary process is constrained in some way so that the daisy population cannot evolve to prefer the environment as it is. Selective advantage is gained by improving the fit between daisy and environment; this can be achieved by changing the daisy or by changing the environment, and evolution will generally opt for the easiest method available. Different factors will affect which method is the easiest, such as the types of daisy in neighbouring patches and their effect on the local environment, and constraints on the range of permissible mutations.

Constraints on evolution are an inevitable feature of any real-world biological system, due to the existence of physical and chemical laws that no system may violate. Chemical laws constrain metabolism, the rate of which typically depends on a number of parameters as some bell-shaped curve. This idea is captured simply in Daisyworld as a growth

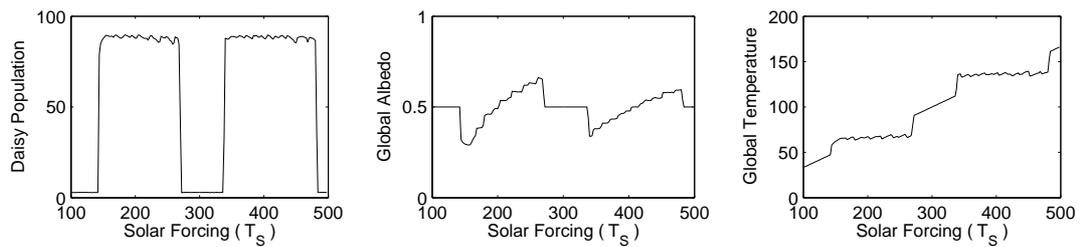


Figure 10.7: Daisy population, global albedo and global temperature for a world where there are two well-separated growth functions (centres at $T = 65$ and $T = 135$) and where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$). Regulation occurs first around one growth function, then around the other.

function that depends on temperature by a Gaussian function, and in the current model by a piecewise linear hat function. While evolution cannot alter the chemical reactions involved in metabolism, it may tinker with the conditions under which those reactions operate to maximise their rate and efficiency, or it may select between different sets of reactions, that is, between different types of metabolism. However, once a metabolism has been chosen during the course of evolution it may often be easier to regulate the environment to suit this metabolism than to switch to a new metabolism entirely. Also, genetic constraints may prevent ‘perfect’ phenotypic adaptation in the short term [74], and while this effect may be lost in long term evolution, it nonetheless creates a potential benefit to regulation.

Different metabolic types may be more successful at different ranges of an environmental variable. In our simplified Daisyworld model, consider a situation where there are two growth functions with centres at different temperatures. The different growth functions may be well-separated, leading to independent regulatory epochs (Figure 10.7), or have overlapping ranges, leading to competitive exclusion (Figures 10.8 and 10.9). When ranges overlap there will usually be one dominant metabolic type around which the environment is regulated, with a flip from one to the other at some critical level of solar forcing. The level at which this occurs depends on the history of the system. Whichever metabolic type becomes abundant first will stop the late-comer from getting a foot-hold in the ecology by holding temperature close to its own optimal level, and thus delay the onset of an ecology (and regulation) based around the other type. This is demonstrated by Figures 10.8 and 10.9 which show competition between two growth functions with overlapping ranges in the face of increasing and decreasing solar forcing respectively (i.e., time flows to the right in Figure 10.8 and to the left in Figure 10.9, although forcing is plotted increasing left-to-right in both).

Another way in which evolution may be constrained and create an opportunity for

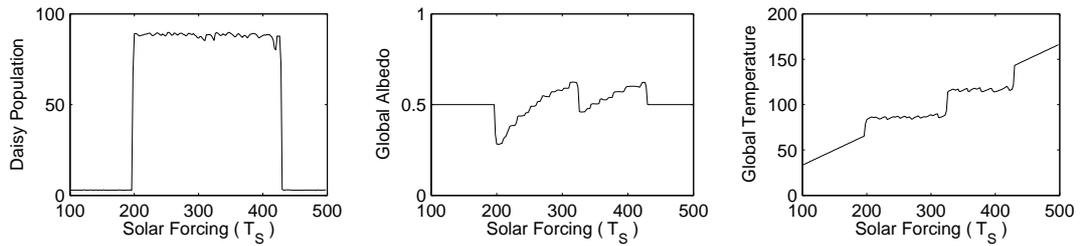


Figure 10.8: Daisy population, global albedo and global temperature for a world where there are two growth functions with overlapping ranges (centres at $T = 85$ and $T = 115$) and where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$). Solar forcing increases over time (time increases from left to right in the plots). Regulation occurs first around the lower temperature growth function, then around the higher temperature growth function; competitive exclusion delays the switch between the two.

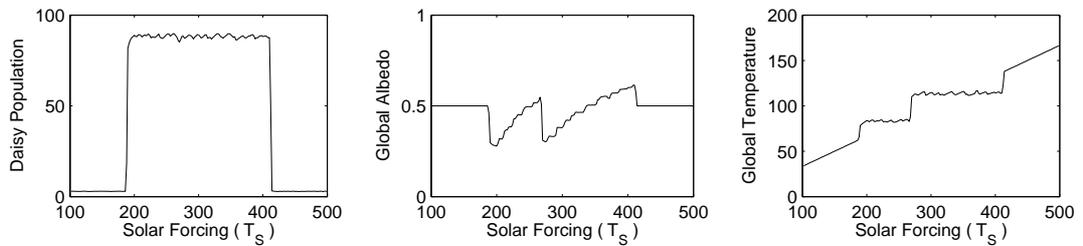


Figure 10.9: Daisy population, global albedo and global temperature for a world where there are two growth functions with overlapping ranges (centres at $T = 85$ and $T = 115$) and where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$). Solar forcing decreases over time (time increases from right to left in the plots). Regulation occurs first around the higher temperature growth function, then around the lower temperature growth function; competitive exclusion delays the switch between the two.

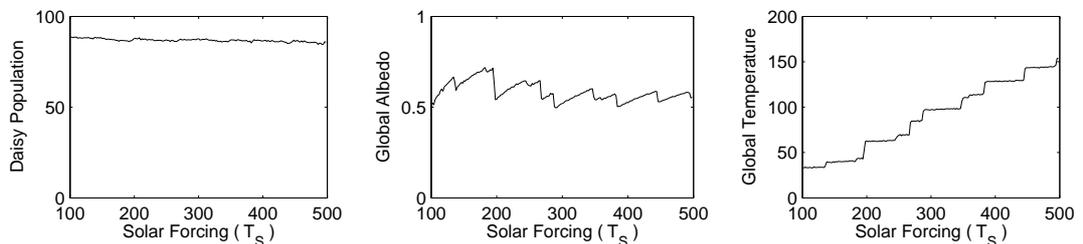


Figure 10.10: Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$) with a probability of 0.2 at each reproduction and where daisy growth function can mutate freely with a probability of 0.002 at each reproduction. The slow mutation rate of preferred growth temperature compared to albedo results in regulatory epochs and stepped increase in temperature.

regulation to evolve is if evolution operates at different rates on different phenotypic traits. Consider the case where the daisy growth function is free to mutate so that it can operate at any temperature and where daisy albedo may also mutate freely to any level between those for black and for white daisies. If both types of mutation occur at the same rate, then the growth function simply tracks the increasing solar forcing and regulation is lost (Figure 10.4). However, if the mutation rate for the growth function is very slow compared to mutation rate of albedo the differential creates an opportunity for regulation. It is easier for a daisy species to evolve a new albedo than a new growth function. This can be observed in Figure 10.10, in which the world is started with a viable daisy population that is then allowed to mutate. At each daisy reproduction (each colonisation of an empty patch), the daisy species may mutate its growth function with probability 0.002 and its albedo with probability 0.2 (so albedo mutates two orders of magnitude faster than the growth function). As can be seen from Figure 10.10, this results in regulatory epochs where the daisy population regulates the global temperature around the optimum for some growth function. Eventually the albedo can mutate no further and mutants with a more suitable growth function can out-compete the existing population to become established as the new dominant metabolic type around which regulation occurs.

10.6 Results 3: Relaxing core assumptions

At the end of Chapter 9 it was noted that Daisyworld is based on two key assumptions:

1. *Organisms create a local buffer against the global environment.* Daisies are assumed to be in patches large enough to maintain a local temperature that is different to the global temperature. It is thus possible for daisy albedo to afford them a selective advantage by altering the local temperature towards the optimum level for growth.
2. *Beneficial adaptations contribute to global regulation.* Selective advantage can only be created by altering local temperature towards a global optimum level for growth. This optimal level becomes the set point around which regulation occurs; thus locally beneficial adaptations are those which contribute to global regulation. There is no possibility of selfish adaptation that destroys regulation.

In this section we will look at what happens when these assumptions are relaxed. First of all we will observe system behaviour when heat flow between daisy patches is assumed to be instantaneous, so that all patches have the same temperature. Daisy albedo still has an effect, but here we assume that any temperature gradients that are created across patch

boundaries are instantly equalised with their surroundings. Thus it is no longer possible for daisies to create a local buffer zone against the global environment.

Figure 10.11 shows the results from a typical simulation run where daisy albedo is allowed to mutate freely and where there is a single universal preferred growth temperature, but where any differences in patch temperature are equalised *at each timestep* by moving all patch temperatures to the global mean. When similar runs were performed with heterogeneous patch temperatures allowed the regulation observed was the most efficient of any of the variations presented here (Figure 10.3). However, in Figure 10.11 we see that no significant regulation occurs when patch temperatures are instantaneously equalised. Regulation does not occur because daisy albedo has no effect on local temperature and all daisies share the same homogeneous environmental conditions. Thus mutation of albedo can offer no selective advantage and albedo has a neutral effect on selection. This means that albedo is a free variable and the daisy population only succeeds when the ambient temperature of the planet happens to be close the preferred level for growth.

When the level of solar insolation allows daisy growth there is a very faint level of regulation. Although all patches are the same temperature and individual daisy species cannot gain a selective advantage, there is still the possibility that the cumulative effects of all the daisies can fortuitously move mean global albedo in the right direction to promote growth. For example, if there happen to be more dark-coloured daisies than light-coloured daisies the global albedo will be below that of bare earth. In a cold environment (low solar luminosity) this will stimulate growth and the daisy population will expand. This expansion of a predominantly low-albedo daisy population will amplify the lowering of the global mean albedo and thus create a positive feedback loop that moves global temperature towards the optimum level for growth. However, this potentially regulatory feedback is brittle, because there is no selective pressure to maintain it. Non-contributing daisies are just as likely to profit from temperature change as those daisies that contribute, allowing genetic drift to dilute the low-albedo population and halt growth. Global temperature will then return to a level similar to that of a dead planet.

The reason that there appears to be some weak regulation of temperature when the daisy population is large is that the effect described above can only occur in the direction of the optimum, since if global albedo moves in the direction to make conditions less suitable for growth the daisy population collapses and the effect is stopped. There is no extension to the range of luminosity in which daisies can grow because there is no selection pressure to coordinate the population; an all-black or all-white population could allow growth at more extreme levels of luminosity but is statistically very unlikely to occur.

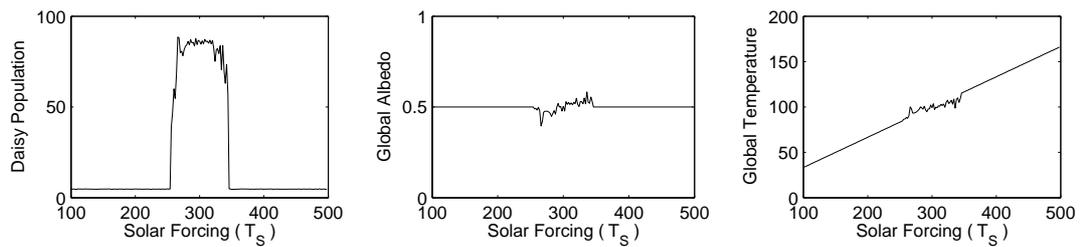


Figure 10.11: Daisy population, global albedo and global temperature for a world where patch temperature instantaneously moves to the global mean at each timestep. There is a single growth function (centre $T = 100$) and where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$). Solar forcing increases over time (time increases from left to right in the plots). No significant regulation occurs because daisies cannot create a local buffer against the global environment in which to create selective advantage by improving their local environment.

Now we will relax the assumption that locally beneficial adaptations always contribute to global regulation. There are two ways in which this can happen: local adaptation can have a neutral effect on regulation, or local adaptation can have a negative effect on regulation. To explore what will happen we need to alter the Daisyworld model so that it is possible for daisies to adapt in ways which provide a selfish benefit but remove or reduce the contribution to global regulation.

Consider the case when local adaptation has a neutral effect on regulation. One way in which this might happen is if selection occurs on organismic traits other than the traits that alter the environment; for example, natural selection might select for daisies with large leaves. In this situation the environment-altering trait becomes selectively neutral. We can see what might happen in this kind of situation by looking at a Daisyworld where colonisation occurs at random and patch temperature is irrelevant to selection. While this is not strictly the same as selection occurring on a different trait, it is a simple proxy for that kind of situation and will serve to illustrate the point.

Figure 10.12 shows the results from a typical simulation run where albedo mutates freely and where empty patches are colonised by random selection of a living daisy species in a neighbouring patch. The daisy population is always large, since solar insolation and planet temperature are irrelevant to growth. No regulation is observed and albedo is a free variable. Local patch temperature reflects the free variation in albedo, and this causes noise in the global temperature.

The results shown in this section show that the Daisyworld assumptions identified above are key to the occurrence of regulation. If daisies cannot create a local buffer via their albedo and thus create a selective advantage, regulation is lost. Regulation is also lost if the link between local adaptive benefit and global regulation is broken, so that adaptation of the environment-altering trait offers no selective advantage.

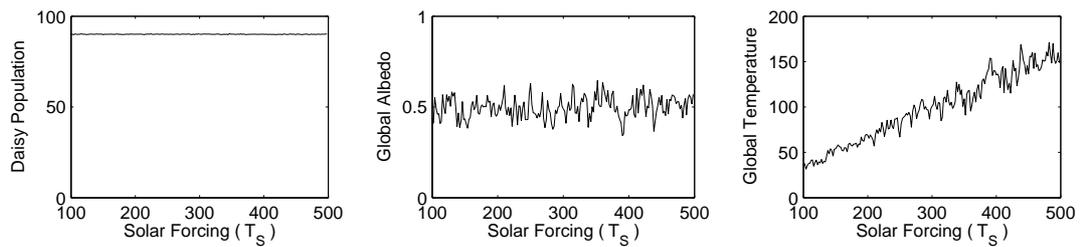


Figure 10.12: Daisy population, global albedo and global temperature for a world where all daisies have an equal growth rate irrespective of patch temperature. There is a single growth function (centre $T = 100$) and where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$). Solar forcing increases over time (time increases from left to right in the plots). No regulation occurs because selection is random and there is no link between local adaptation and global regulation.

10.7 Conclusion

We have presented a model that is derived from Daisyworld, but is simplified and extended to allow for a more comprehensive study of the compatibility of biotic environmental regulation with evolutionary theory. The model has been described here using the language of Daisyworld (daisies using albedo to regulate temperature in the face of solar forcing) but the mathematical formulation of the model is actually very general, meaning that similarly constructed models may be used to study biotic regulation in other scenarios. Our model shows that what is needed for regulation to emerge are constraints on the evolutionary process and the possibility for organisms to create a local buffer against the global environment, criteria that we feel are plausible in a wide variety of biological systems.

Relaxing the assumptions of the model shows that two Daisyworld assumptions (that organisms can create a local buffer against the global environment, and that local adaptation contributes to global regulation) are key to the evolution of regulation. If either of these assumptions is relaxed then regulation does not occur. However, varying some of the physics of the model appears to have no effect on the evolution of regulation. Here we used a simple piecewise linear growth function rather than the parabolic growth function of the original Watson and Lovelock model, but regulation was still observed. Heat radiation dynamics were also much simplified, without any effect on the qualitative nature of the results achieved.

It appears from this work that the key features of the Daisyworld model which are necessary for regulation to evolve are a peaked growth function with an optimal response at some fixed level of the environmental factor, a local buffer allowing organisms to beneficially affect their immediate environment and thus gain selective advantage, and constraints on the evolutionary process so that there is a need for regulation. These criteria

all seem likely to be satisfied in some situations in the natural world. Chemical reactions (including metabolic reactions) typically have a peaked response with a maximum rate achieved in ideal conditions. Constraints on evolutionary adaptation certainly exist, coming from physical and chemical laws as well as from evolutionary canalisation and genetic constraints on phenotypic adaptation. Perhaps the most contentious criterion is the existence of a local buffer within which an organism can selfishly benefit from improving its local environment, but this criterion may be less important in the real world than it is in Daisyworld.

It is easy to think of some situations where an organism will create a local buffer against the global environment; one example that is particularly relevant to Daisyworld is the boreal forests, where the lower albedo of the trees compared to the surrounding tundra permits rapid warming and an earlier start to growth in springtime than would otherwise have been possible [104]. However, there are also many situations where the possibility of a local buffer is harder to imagine, but where regulation still occurs. For instance, chemical regulation of the oceans by the marine biota seems to happen despite the situation of the marine organisms in a liquid matrix which would rapidly disperse any excreted materials that might contribute to the formation of a favourable buffer zone. The evolution of environmental regulation in situations where it is hard or impossible to form a local buffer will be examined using the Flask model in following chapters.

Chapter 11

The Flask model

11.1 Overview

In this chapter we introduce the Flask model, a new and original model intended to explore the evolution of environmental control by the biota in a scenario where some of the key assumptions of Daisyworld are relaxed. The Flask model is based on evolving populations of microbes that affect their environment as a side effect of metabolism. It is hoped that the Flask model will be able to address some of the criticisms applied to Daisyworld and to contribute fresh ideas to the debate over Gaia theory. However, we do not intend that the Flask model should be limited to Gaia theory; it is designed as a general model of biotic environmental control that can be applied on a variety of different spatial and temporal scales, of which Gaia is just one. For instance, at a smaller scale the Flask model could be used as a model of niche construction, where organisms change the selection pressure they experience by altering their environment.

The next section will recap on some of the arguments presented in previous chapters and motivate the development of the Flask model. This will be followed by a conceptual overview of the Flask model in which its key assumptions will be identified together with some areas where preliminary study is needed before the main hypothesis can be approached. Due to time and space constraints, the Flask model is not explored to its full potential in this thesis and while some preliminary experimental work is presented in the following chapter, a large amount of work is left for future research.

11.2 Moving on from Daisyworld

Ecological models often ignore the interaction between the biota and the environment. For example, Lotka-Volterra population models [115, 179] and replicator equation models [83] typically look only at the interaction between species. The external environment, where it is modelled at all, is usually taken only as a (possibly spatial) source of food; other aspects that may affect the growth and reproductive success of organisms are ignored.

However, it is increasingly recognised that there is a constant and significant interaction between living things and the abiotic environment. Ideas of niche construction [102, 103] and the extended phenotype [40] highlight the importance of the environment in evolutionary dynamics, while the increasingly accepted Gaia theory [116] suggests that not only do organisms affect their environment, they do so in a way that regulates the biosphere to conditions that are suitable for life. In recent years, artificial life models of ecologies have begun to include coevolution between species; it is now important that practitioners recognise the fundamental importance of coevolution between the biota and the abiotic environment.

We are interested in the possibility of regulation of the abiotic environment by the biota, which on a planetary scale is called Gaia theory but need not be confined to this scale. The interaction between life and its environment is a fascinating topic at any scale, from microbial mats involved in the formation of stromatolites [141], to beavers that act as ‘ecosystem engineers’ creating lakes and canals [191], to cloud formation caused by marine phytoplankton [27].

In the previous chapter we looked at a version of the Daisyworld model [183] in which the physics had been made simpler than those of the original in order to clarify the basic properties of the model. We looked specifically at the role of evolutionary adaptation in the Daisyworld model, and examined the conditions under which environmental regulation could evolve. It was found that two basic assumptions of the Daisyworld model are essential for regulation to emerge. The first of these is that organisms should be able to create a local buffer against the global environment, within which to gain selective advantage by favourably affecting their immediate locale. The second assumption is that locally advantageous adaptations must also contribute to global regulation, that is, that selfish adaptations change the environment in the direction necessary for global regulation to occur. We also saw that regulation only emerged when there was a need and an ability to regulate. If daisies could adapt their preferred growth temperature to match the current global temperature then there was no need for regulation, but if constraints were placed on adaptation of growth temperature then it became beneficial to moderate local

temperature via an appropriate albedo.

Armed with our new understanding of Daisyworld dynamics, it is now time to move on. The Daisyworld models have been the source of many interesting debates and have contributed greatly to the debate over Gaia theory. Whatever people may think about the Gaia hypothesis, they are forced to concede that Daisyworld does at least provide a cybernetic example of how regulation might emerge from the dynamics of a non-teleological model based on sound ecological and evolutionary principles. However, it seems likely that the law of diminishing returns will apply and that the number of new insights that will result from new variations of Daisyworld will be few. There is a recognised need for new models [177].

Here we present the first steps towards a new model, the Flask model, which we hope will answer significant questions about the logic of environmental regulation by the biota. The Flask model is conceptually simple, but has the scope for a wide variety of questions to be asked within its framework. It consists of looking at the ecological and evolutionary dynamics of microbial populations in a series of microcosms, which may be thought of as flasks. Each flask contains a well-mixed solution of nutrients, chemicals and neutral substrate liquid. Adding microbes to the flask creates conditions in which communities may form and evolutionary change occur. By enforcing perfect mixing *within* flasks, but allowing only partial mixing *between* flasks, we create the conditions for multi-level selection. Between-flask selection takes place at the ecosystem level and within-flask selection takes place at the level of the individual microbe.

Due to constraints regarding both time and space, the results presented in this thesis are preliminary to a more in depth future study of the Flask model. While the eventual aim is to study the evolution of environmental regulation by the biota, we must be content here with taking only the first steps towards this goal. In this chapter we present the model and argue why we might expect to observe the evolution of environmental regulation within it, and also why it might be a good model for the evolution of environmental regulation in nature. We identify several basic criteria that must be satisfied before the model can be used to test any larger hypotheses concerning the evolution of environmental regulation, which include the demonstration of heritability at the community level in the Flask model, demonstration also of a correlation between environmental regulation and biomass within individual flasks, and confirmation that the spatial structure of the model does indeed enable the conditions for multi-level selection.

11.3 Biological basis for the Flask model

The basic Flask concept of microbial communities collectively altering their shared environment is in part inspired by the artificial ecosystem selection experiments performed by Swenson et al [162, 163]. In these experiments, self-contained microbial ecosystems were created in test tubes under laboratory conditions. After a fixed development period, the chemical composition of the liquid in each test tube was measured for pH. In one set of experiments, the ecosystem displaying the highest pH was used to provide seed microbes for a new generation of test tube ecosystems, while in another series of experiments the lowest pH community was chosen. A sustained response to selection was observed, with deviation of pH away from the starting level increasing at each generation.

The experiments performed by Swenson et al show that selection on microbial communities for their effect on their environment can over time produce communities that alter their environment in a particular way. This suggests that if a particular type of environment increases microbial growth, community-level selection might lead to communities that move their environment closer to optimal growth conditions. This effect could result in environmental regulation.

The Flask system is modelled as an array of connected flasks, each containing a microbial population suspended in a liquid medium. The liquid medium is important as it allows the assumption of a near-homogeneous environment inside every (well mixed) flask. Microbes are chosen as the model for the simulated organisms for a number of reasons. They have simpler metabolisms and foraging techniques than multicellular organisms, and are able to take advantage of chemical nutrients absorbed in liquid form from the environment. Microbes are also capable of using a wider variety of metabolic reactions than the evolutionarily canalised metabolisms of multicellular organisms, making the possibility of evolving to use a different metabolic reaction with just a few mutations less implausible. The microbial reproductive cycle is also relatively simple, being asexual in nature and producing fully functioning offspring.

Microbes are also known to coexist in tightly integrated symbiotic communities called microbial mats that can be found all over the Earth's surface and in a wide variety of different environments [126]. Microbes arose much earlier in the history of life on Earth than multicellular organisms, which means they are much more likely than multicellular organisms to have played a part in regulating the biosphere; after all, multicellular organisms could never have arisen if there hadn't been a suitable environment for them to evolve in. The wide global coverage and geophysical action of microbial species also supports the idea that microbes could play a significant role in regulating the global environment.

There are a variety of real environments to which the Flask model might be com-

pared, which basically include any liquid environment populated by microorganisms. The precise suitability of Flask for representing a particular environment will be in part determined by the rate of mixing between flasks. While mixing within a flask is assumed to be perfect, the level of between-flask mixing can be tweaked to fit different scenarios. If between flask mixing occurs at quite a high rate, the model could represent an ocean environment, with individual flasks representing different water cells. If the rate of between-flask mixing was lower, then the Flask model might be compared to the bacterial communities in mud flats or marshes. Depending on how the between-flask mixing is modelled, the Flask model could also be seen as a form of island model, or perhaps an arrangement of inter-tidal rockpools.

However, despite this range of possible analogies, the fact remains that the Flask model is primarily of heuristic value in determining the logic of environmental regulation. If Flask can be used to test the implications of various assumptions, or to find a theoretically sound way in which environmental regulation can evolve in accord with Darwinian theory, then it will have been vindicated.

11.4 Differences from existing models

The Flask model relaxes the two main assumptions of the Daisyworld model. In Daisyworld, there is a well-defined local buffer within which daisies can create selective advantage by ameliorating the impact of the global environment. In Flask, this buffer is expanded so that individuals cannot by themselves create a selective advantage by altering the environment. Such an advantage can only exist through the collective result of the environment-altering traits of all individuals inside a flask. The only buffer that remains in the Flask model is that each flask maintains a homogeneous internal environment that is near-closed to input from other flasks.

Also, in Daisyworld, adaptations which create local advantage are also those which contribute to global regulation. There is no possibility for a daisy species to succeed by adapting in a way that would cause global regulation to break down. In the Flask model however, individual level mutations to the environment-altering trait are selectively neutral, since all individuals in a flask share the same environment. Individual-level selection acts on metabolic traits, and it is quite possible for a beneficial adaptation to metabolic function to occur in conjunction with a mutation to the environment-altering trait that pushes the environment further from the optimal conditions for growth. Selfish adaptations can thus result in regulatory, non-regulatory, or neutral effects on the abiotic environment.

Another advantage of the Flask model over the Daisyworld model is that it incorpo-

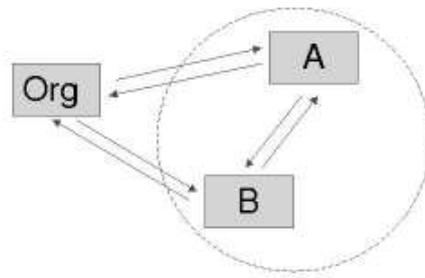


Figure 11.1: Every organism has an environment (dotted circle) which consists of both abiotic (A) and biotic (B) components. There is two-way feedback between the organism and all of these components, as well as between the components and each other.

rates a more realistic ecological model, including the formation of food webs and nutrient recycling loops. The same is true of the Guild model, but the Guild model shares the same basic assumptions as the Daisyworld model, the limitations of which are described in previous paragraphs.

11.5 The Flask model: A conceptual sketch

Every organism has an environment, which is made up of both biotic components (e.g., other organisms) and abiotic components (e.g., temperature, chemical fluxes, geography, salinity, etc.). The organism interacts with all of these components, which also interact with each other (see Figure 11.1). Since organisms evolve, it is possible that selection may occur on the environment-altering traits of organisms, so that traits are favoured which change the environment in a way beneficial to the organism. Since many organisms share similar basic preferences for some aspects of the environment (such as temperature, pH, etc.), the global result of many selfish adaptations may be environmental regulation.

The basic unit in the Flask model is a flask containing a neutral liquid matrix (see Figure 11.2). The flask receives a chemical influx at a fixed rate and also loses stored chemicals at a fixed rate through chemical outflux. Thus the chemical composition of the liquid medium within the flask will fall to a steady state equilibrium in the absence of perturbation. A microbial population is introduced to the flask, which may contain individuals capable of metabolising some of the chemicals stored in the flask and thus create the conditions for a microbial ecology to form.

Microbial metabolism involves the consumption of chemical nutrients in a genetically determined ratio specific to the individual. Some of the consumed nutrients are converted to biomass, with the remainder excreted to the environment as chemical waste (also in

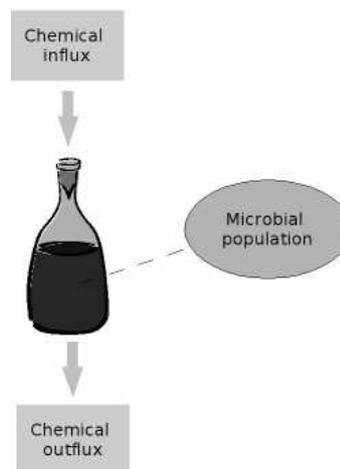


Figure 11.2: Schematic of a single flask. The flask receives a steady chemical influx and loses a fixed proportion of stored chemicals by chemical outflux. A microbial population is introduced to the flask and affects its chemical composition.

a genetically determined ratio). One microbe's chemical waste may potentially be consumed by another microbe species.

Microbes reproduce asexually by splitting when they reach a certain biomass threshold, with a low probability of mutation. Microbes also lose biomass at a steady rate due to the cost of living, assumed to be lost to the environment as non-recoverable heat or energy. Microbes die if their biomass falls below a certain level.

Within each flask there are all the ingredients necessary for an ecology to form. There is a steady supply of nutrients to the environment. Microbes can consume and excrete different chemical products, so it is possible for one species' excreta to be food for a different species. Reproduction or death depend on the fit between a microbe's metabolism and the food available, which creates selection pressure on individuals since metabolism is genetically specified.

Not all of the chemicals in the flask are used in microbial metabolism. Some chemicals are consumable, while others are non-consumable chemicals that form part of the abiotic environment of the microbes. Although the non-consumable chemicals are not used in the metabolic process, it is assumed that microbial activity has an effect on their concentration within the flask. This takes the form of a genetically determined effect on the level of each non-consumable chemical, that may be implemented as a side-effect of metabolism or as a fixed rate contribution from each microbe. The microbe population as a whole interacts with the input and output fluxes of the non-consumable chemicals to determine the composition of the abiotic environment.

In addition, the composition of the abiotic environment affects the rate at which mi-

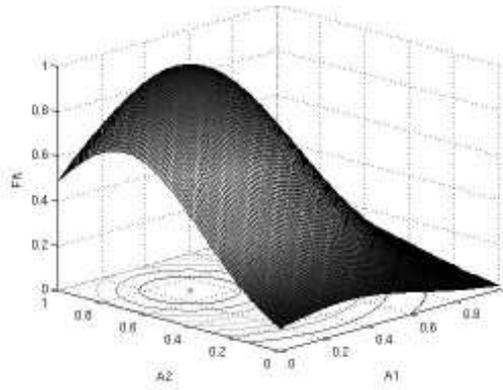


Figure 11.3: Metabolic rate is affected by the composition of the abiotic environment. Plot shows an example function giving metabolic rate as a function of the composition of a two-component abiotic environment. There is assumed to be an ideal abiotic composition, at which the environment fits metabolism perfectly, which is shared by all organisms.

crobes can metabolise nutrients (as, for example, the environmental pH level might affect the growth of real bacteria). It is furthermore assumed that there is a single preferred state vector for the composition of the abiotic environment, at which the metabolic rate of all microbes is maximised. An example function is given in Figure 11.3 for a two-component abiotic environment.

The genetic code of a microbe specifies the ratio in which it consumes nutrients, the ratio in which it excretes waste, and its effect on the abiotic environment. Since the flask is assumed to be well-mixed, the environment is experienced equally by all microbes. This affects the ways in which selection pressure can affect individuals. Selection can clearly act on phenotypic traits concerned with metabolism, since the availability of suitable food will determine growth and hence reproductive success. However, there is no possibility of individual-level selection on phenotypic traits which affect the abiotic environment, since the well-mixed liquid medium in a flask means that the environment is experienced equally by all microbes within it. Thus there can be no individual benefit arising from an environment-altering trait, meaning that traits affecting the abiotic environment are selectively neutral at the level of the individual. In the absence of any higher-level selection, this makes the collective effect of the microbial population on the abiotic environment a free variable. Even though the composition of the abiotic environment will have a large impact on the overall success of the microbial population, there is no way that it can be beneficially steered by evolutionary adaptation within a flask.

Now, consider the case where there is an array of many flasks (see Figure 11.4). If each of the flasks is initialised with a different microbial population, the different starting

conditions should cause the ecological and evolutionary dynamics within each flask to create different ecological and chemical conditions. In some flasks, the population will quickly die out and the chemical environment will return to its abiotic stable state. In other flasks, viable communities will be formed with functioning food webs.

The success of a community (measured in growth rate or number of individuals) will depend on the interaction of the microbial population with its environment. A community which only consumes the nutrient influx and does not evolve good nutrient cycling loops will be limited by the size of the nutrient influx. By contrast, communities which form stable food webs will recycle nutrients many times before they leave the system via nutrient outflux, and will thus make available a much larger food supply that can sustain a larger microbial population [55, 176]. In order for the model to be thermodynamically correct, either some energy must be lost during each metabolic reaction (to prevent infinite recycling), or an additional assumption must be made of an external energy source such as sunlight, making the Flask model an energetically open system rather than a closed system.

The interaction of a community with its abiotic environment is of key importance. Since the abiotic environment affects the metabolic rate of the microbes, the collective influence of the community will result in an abiotic environment that favours growth to a greater or lesser degree. We should therefore be able to observe a correlation between community biomass and the distance of the current abiotic environment composition vector from its ideal.

Note that it is extremely unlikely that a single species could maintain optimal environmental conditions alone, since it will only be able to push each environmental factor in one direction and would need to alter each factor by a precise amount. It is theoretically possible that a flask community consisting of a single species with exactly the right kind of genotype could stabilise at a size that optimised the environmental factors in balance with the fluxes through the flask, but the chance of this happening in any non-trivial system are vanishingly small. Also, if this event did occur, it would be very brittle to mutation. This means that the creation of optimal environmental conditions must be a collective activity of a stable food web, and it is infeasible for a single species to do so alone.

Different communities will have ecologies that are more or less successful, as determined by their efficiency of nutrient recycling and their effect on the abiotic environment. These differences will be expressed as variations in biomass or the number of individuals, and these quantities could form the basis of some selection criteria at the level of the community. However, what is not clear is how these communities could ever ‘reproduce’ in order for selection pressure to have an effect.

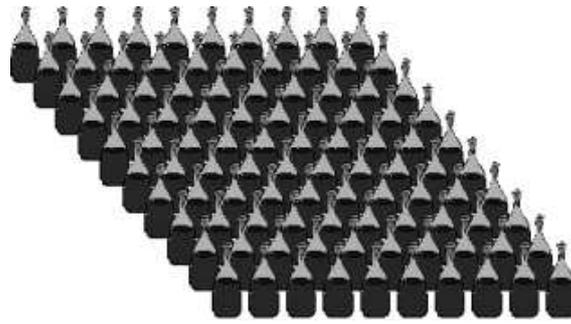


Figure 11.4: Ecological and evolutionary dynamics mean that microbial communities within flasks may be more or less successful depending on initial conditions. If migration is possible between flasks this raises the possibility of colonisation or invasion by more successful communities.

Consider the case where each flask in the array is connected to its neighbours by a valve that allows for limited mixing of the liquid matrix between flasks, and thus enables migration of microbes from one flask to another. With the possibility of migration between flasks, it might be expected that viable communities might spread from populated flasks to lifeless ones. What is more, since the populations in different viable flask communities might be of different sizes, some communities might be more or less successful at colonisation than others. We might expect more populous flask communities to be better colonisers since they can spawn greater numbers of potential colonist individuals. This might lead to more rapid expansion by more populous communities, and perhaps invasion of smaller communities by larger ones.

We have now described means by which flask communities can be said to display variation, and also how they might selectively reproduce. The final ingredient for some form of community-level selection to occur is heritability. Given that the flask communities are made up of individual microbes that are subject to mutation, should we expect a colonist community to develop similarly to its parent community? It seems likely that the answer to this question depends on timescale.

In the short term, community dynamics are dominated by ecological interactions, with different species reproducing dependent on food supply and competition. If the mutation rate of individuals is sufficiently low, we might expect a reasonable degree of constancy in the community species distribution over short time periods. However, over longer period mutations and evolutionary dynamics come into play and there is less reason to expect any constancy. New species may arise that have a significant effect on community composition. So, we might expect to observe similarity between colonist communities and their parent communities in the short term, but not in the long term. This issue will be developed in the following chapter with experimental results concerning the time

dependent nature of ecosystem heritability.

The microbial communities within a flask must make good use of their nutrient supply to be viable, and differential success between alternative viable communities will be in large part determined by their effect on their abiotic environment. Since the success of a flask community will affect its likelihood of colonising or invading other flasks, we now have a possible criteria for selection between flasks. Over time we might expect to see the flask array dominated by communities that keep the composition of their abiotic environment close to the ideal levels for growth.

Group selection has historically been thought to be a weak or non-existent force in evolutionary theory, but recent studies have shown that there are specific circumstances where it can play a significant role. These circumstances are those where groups are largely isolated but occasionally mixed, so that individuals from successful groups come to form a large proportion of the entire population [156].

These conditions can be met in the Flask model dependent on the rate of migration between flasks. If flasks are completely isolated no migration is possible, and while some flasks may contain more successful communities than others, there is no possibility of colonisation or invasion, and selection cannot occur. On the other hand, if flasks are well connected and there is no barrier to migration, there is no community isolation and group benefits cannot occur within flasks. However, it seems plausible that if the rate of migration is present but small, then conditions for group selection might be met and the possibility of selection acting on different flask communities might exist.

The main hypothesis that the Flask model is designed to test is given below:

If there is feedback between organisms and their environment, so that organisms affect their environment and growth is maximised in a particular set of environmental conditions, and if there is an appropriate spatial structure that allows partial isolation of semi-closed communities, then multi-level selection will result in the evolution of communities which collectively regulate their environment to conditions suitable for life, in a manner consistent with neo-Darwinian evolutionary theory.

The Flask model so described rests on a large number of conjectures and assumptions, and the author does not expect the reader to accept these without evidence and clarification. A quantity of work is needed to show the validity of these assumptions; in the next chapter we will present the first steps in this process. However, we feel that while the details of the model need proof and refinement, the model is logically consistent both in itself and with modern evolutionary biology.

11.6 Basic assumptions and points to prove

In this section we discuss some assumptions of the Flask model that require clarification and supporting evidence if the Flask model is to be valid

11.6.1 Universal preferred environment?

The existence of a universally preferred set of environmental conditions for growth is a simplifying assumption of the Flask model. As discussed in Chapter 10, physical and chemical constraints exist that restrict the range of conditions in which biological growth can occur. Typically, the rate of a given metabolic reaction varies with respect to a particular environmental factor as a bell-shaped curve (as it does in Flask), and this curve cannot be altered by biological adaptation. However, the choice of which of the many possible metabolic reactions is used to extract energy from the environment *is* subject to adaptation, and given that different reactions may have different optimal conditions in which their rate is maximised, the Flask assumption of a universally preferred environment may seem to be flawed. However, we argue here that subject to certain caveats, this is not the case.

While different metabolic reactions may have different optimal conditions, these conditions are often broadly similar. It seems likely that adaptive choices made far back in evolutionary history mean that the differences between different metabolisms are less significant than their similarities, and that many metabolic types present in the natural world are variations on a theme. For instance, different varieties of oxygen-based metabolism may have slightly different optimal conditions, but these are all highly similar when compared to any form of methane-based metabolism.

Also, evolutionary canalisation may mean that it is very difficult or impossible to suddenly switch to a new form of metabolism with a radically different set of preferred conditions. Such a change is likely to require an unfeasibly large number of coincident mutations and it is in any case difficult to see how such a mutant could survive long enough for its traits to reach fixation. Timescale is also an issue; while it may be possible to switch to a different form of metabolism with a different set of preferred conditions, the difficulties in doing so might lead to long periods of fixation on a single metabolism type. This might give rise to a form of homeorhetic regulation, in which regulatory epochs occur around different metabolic types in sequence.

We feel that these reasons justify the simplifying assumption that all organisms in the Flask model have a universal preference for a particular set of environmental conditions. Nonetheless, the implications of multiple preferences raises intriguing theoretical

questions that may be explored by extension to the Flask model in future work.

11.6.2 Environmental fit correlates with biomass?

It is reasonable to suggest that when the abiotic environment has an effect on microbe growth rate, so that growth is maximised at some particular environmental composition, there will be a correlation between the size of the microbial community in a flask and the deviation of the actual environmental conditions from the optimum. Real world examples might include the relation between temperature and growth in plants, or between bacterial growth and the pH level of their surrounding environment. However, this claim must be verified in the Flask model before any larger hypotheses can be approached.

11.6.3 Heritability at flask level?

For any higher-level effects to give selection for communities that regulate their abiotic environment, it is first of all necessary to show that there is some measure of heritability between a parent flask community and an offspring colonist community. If there is no heritability, so that colonist communities bear no relation to their parent communities, there is no reason to expect any persistence or spreading of communities that have a favourable effect on their environment. This would preclude any higher-level selection that might lead to the evolution of communities that regulate their environment.

11.6.4 Spatial structure leads to group selection?

Even if heritability can be demonstrated at the level of flask communities, it remains to be seen whether or not the spatial structure of the Flask scheme, with perfect intra-flask mixing but weak inter-flask mixing, gives rise to any form of group selection dynamics. It has been shown that group selection can be significant if the groups are largely isolated with occasional mixing [156]. It seems likely that the spatial structure of the Flask model will provide these conditions and allow higher-level selection to take place, but this is far from inevitable and must be demonstrated.

11.7 Conclusion

In this chapter we have presented a sketch of the Flask model, and described the work that must be performed before the main hypothesis for which it is designed can be tested. Time and space constraints on this thesis mean that the majority of this work is left for

the future, but in the next chapter we will address the issue of whether or not heritability can be demonstrated at the community level in the Flask model.

Chapter 12

Response to ecosystem selection in the Flask model

12.1 Overview

In this chapter we describe an experiment designed to show whether or not community-level heritability can be demonstrated in the Flask model. Using a simplified version of the model, that treats all flask communities as completely isolated and ignores the feedback from the abiotic environment to the growth rate of the microbial population, we perform tests to measure a response to selection at the ecosystem level. A population of different flask communities is initialised and allowed to develop for a fixed period. At the end of this period a fitness score is assigned to each flask by measuring the error between the actual composition of the abiotic environment in each flask compared to a pre-specified target vector. Artificial selection is then used to create a new population of flask communities, performed by inoculating sterile flasks with samples of the microbial communities from the flasks with the lowest error. Results from this experiment show that over time a response to artificial ecosystem selection is observed in the Flask model.

We begin with some background to the debate in evolutionary theory over group selection and related modelling work in this area. This is followed by a description of the methods used and results obtained, before some discussion of the implications of these results for the Flask model and for evolutionary theory.

12.2 Background

In the previous chapter we outlined some criteria that must be met if the Flask model is to be viable. One of these criteria is that there must be some heritability at the level of the flask communities, so that a colonist community resembles its parent community. If this cannot be shown to hold, then there is no possibility that higher-level selection effects can lead to the evolution of flask communities that regulate their abiotic environment. Higher-level selection is a contentious issue in evolutionary theory, but there are grounds to believe that group selection can play a significant role in a scenario such as the Flask model.

Lewontin [113] states that any level of biological organisation that can be grouped into a population of units has the potential for evolution by natural selection. He goes on to describe the three necessary and sufficient properties that are required for evolution to take place: (1) phenotypic variation, (2) differential fitness based on phenotypic variation, and (3) heritability of fitness. Lewontin goes on to argue that selection at higher levels than the individual is unlikely, since higher-level units are unlikely to exhibit heritability of phenotypic traits. This opinion is similar to the gene-centred view of evolution propounded by many evolutionary biologists (e.g., ‘selfish gene’ theory [39]). Group selection has been the subject of various critiques, notably from Maynard Smith [130] and Williams [186], and until recently the idea of selection at a level higher than that of the individual has not been favoured.

However, higher level selection has latterly been a topic of renewed interest, with both theoretical and experimental evidence for its existence being presented [71, 156, 181]. What used to be known as ‘group selection’ is beginning to find new acclaim as ‘multi-level selection’, though it is still viewed with suspicion by many.

Recent work has demonstrated that artificial selection at the level of the ecosystem can lead to a sustained evolutionary response [162, 163]. In a number of laboratory experiments a statistically significant response was observed to artificial selection, both when soil communities were selected for the dry weight of plant biomass supported and when pond-water communities were selected for raising or lowering pH level [162]. In a similar experiment, Swenson et al showed that artificial ecosystem selection could be used to evolve microbial communities to break down the environmental pollutant 3-chloroaniline [163]. The selection in these experiments is artificial, not natural, and it is possible that selection was implicitly for individual-level traits, but the results obtained are significant and encouraging for our current purpose.

Some simulation experiments into ecosystem selection have recently been reported in the artificial life literature [137, 138]. Using a system based on Lotka-Volterra population

equations with mutation between populations, Penn showed a response to selection for diversity [137]. Using a similar system, Penn and Harvey showed a response to selection on species composition ratio in a Lotka-Volterra model without mutation, i.e., where the only variation in ecosystems was caused by changes in species distribution [138].

A possible weakness of Penn's model is that selection is performed directly on characteristics of the population itself, i.e., diversity and species composition. It is possible, though not certain, that this introduces some unwanted bias into the evolutionary process. Whether or not this is true we feel that a clearer understanding could be gained from experiments where selection is for a trait that is affected by, but external to, the population. Selection for the effect of the ecosystem on its environment meets these criteria.

The issue with group selection is not whether it exists, but to what extent. Individual-level selection has traditionally been thought to be the dominant force, and in most situations this is the case. However, there are situations where higher-level selection pressure can play a significant role in shaping evolutionary adaptation and we feel that the Flask model may be one of these. As a precursor to examining higher-level selection in the full Flask model, we are here concerned with demonstrating the possibility of a response to selection at the level of the flask community.

12.3 Method

Here we seek to demonstrate in the Flask model a similar community-level response to selection to that observed in Swenson's pond-water ecosystems [162]. We apply a similar experimental method to Swenson et al, where flask ecosystems are selected for the composition of their abiotic environment in a manner analogous to the selection for pH in Swenson's experiment.

Conceptually, the Flask model is set up as described in the previous chapter, with some simplifications. The flasks are assumed to be isolated from each other and selection is performed artificially, in order to avoid any complications arising from the migration/colonisation/invasion model of community selection. Individual flasks are modelled as described except that there is no feedback from the abiotic environment to microbe growth; again this is done to reduce the complexity of the model and allow us to clearly observe the phenomenon of interest, which is a response to selection.

The model simulates the growth and evolution of microbial communities suspended in a liquid medium. This medium is held in flasks supplied with continuous chemical fluxes. Individual microbes grow and reproduce dependent on food supply, and the ecological interaction between them leads to the formation of food webs in each flask. Mutation may occur during reproduction, allowing the genesis of new microbial strains.

The composition of the liquid medium in each flask determines the environment of the microbes. Some of the chemicals present may be consumed as food by the microbial population and converted to biomass, while others are non-consumable and form part of the abiotic environment. In addition it is assumed that the liquid medium has properties such as temperature, pH, salinity, etc., and that these can be affected by microbial activity. We will refer to these non-consumable chemicals and physical properties of the flask environment collectively as ‘abiotic factors’ for ease of discussion; while chemical nutrients are also abiotic we feel that their role as the subjects of metabolism justifies this notational convenience. The effect of the microbes on abiotic factors is modelled here as a side-effect of metabolism, with a genetically specified effect caused by each microbe for every unit of biomass created.

The composition of the abiotic environment resulting from the interaction of the input and output fluxes with the collective actions of the microbial population forms a ‘phenotypic’ ecosystem trait that is used as the basis for selection. Offspring ecosystems are formed by inoculating sterile flasks with seed populations sampled from the most successful ecosystems in the previous generation. The response to selection is measured as the change over time in the distance of the environmental state variables from some pre-specified ideal ratio.

The experimental method is inspired by the method of Swenson et al [162, 163]. We wish to replicate the results of their *in vitro* experiments in our *in silico* experiments, in order to show that a response to artificial ecosystem-level selection can occur in the Flask model. If such a response can be shown, it implies that there is sufficient heritability between parent and offspring communities for the occurrence of the community-level selection necessary for the evolution of environmental regulation.

12.3.1 Flasks

Each flask contains a neutral liquid matrix in which is suspended a microbial population. There is a flow of liquid through the flask. The inflow brings with it steady influxes of chemicals and steady inputs to abiotic factors, while the outflow removes a fixed proportion of stored chemicals and a steady output from abiotic factors. Flasks are well-mixed, so that in the absence of perturbation the composition of each flask will reach a steady state equilibrium. Each microbe both consumes and excretes chemical nutrients, and also has an effect on the levels of the abiotic factors.

The state of the flask is therefore a vector V of length $M + N$, where M is the number of chemical nutrients and N is the number of abiotic factors. The update to V at each timestep is given by Equation 12.1, where F_{in} is the influx vector, F_{out} is the outflux

vector, and E is the net effect on V of the microbial population.

$$\Delta V = F_{in} + E - F_{out} \quad (12.1)$$

The environmental state vector V can be subdivided into separate vectors for the consumable chemicals and abiotic factors, so that $V = (X, Y) = (x_1, \dots, x_M, y_1, \dots, y_N)$ where X and Y are the consumable and abiotic parts respectively.

12.3.2 Microbes

Microbes are modelled as simple bacterial organisms that consume and excrete chemical nutrients in a genetically specified ratio. Each microbe also has a genetically specified effect on each abiotic factor as a by-product of metabolism. Each microbe can therefore be represented by a vector $(B, \lambda, \mu, \alpha)$ where $\lambda = (\lambda_1, \dots, \lambda_M)$ represents the consumption ratio, $\mu = (\mu_1, \dots, \mu_M)$ the excretion ratio, and $\alpha = (\alpha_1, \dots, \alpha_N)$ the effect on the abiotic environment. Clearly $\sum_i^M \lambda = 1$ and $\sum_i^M \mu = 1$ hold since all materials consumed and excreted must be accounted for; there is no such constraint on α since the effect of the microbe on the abiotic environmental factors does not necessarily involve mass transfer and is thus treated generally. B is the current biomass of the microbe and is a variable, whereas λ , μ and α are genetically fixed for the lifetime of the individual.

12.3.2.1 Genotype

The genotype for a microbe is an array with $2M + N$ loci taking values in the range $[-1.00, 1.00]$. There are two loci for each chemical nutrient, specifying what proportion of the microbe's total consumption and excretion is constituted by that chemical. The M loci for consumption are linearly mapped to the range $[0.00, 1.00]$ and normalised, to give the ratio in which chemicals are consumed; excretion ratios are found similarly. For example, if $M = 3$ and the consumption loci of the genotype are $(-0.4, 0.7, 0.1)$, this would map to $(0.3, 0.85, 0.55)$ and give a normalised consumption ratio of $\lambda = (0.18, 0.5, 0.32)$. There are N loci for effects on abiotic environmental factors, which map directly to the effect the microbe has on each factor for every unit of biomass created.

12.3.2.2 Metabolism

At each timestep, each microbe will try to consume up to C_{max} units of food in the consumption ratio determined by its genotype. If sufficient of each nutrient is present to satisfy this demand, the food is consumed and converted to biomass with a standard efficiency of θ , with the waste being excreted (i.e., 10 units of food consumed with efficiency

of 0.6 makes 6 units of biomass and 4 units of excreta). The excreta is returned to the environment in the ratio determined by the microbe's genotype.

Each microbe incurs a biomass decrement of γ units at each timestep for the cost of living, assumed to be lost to the environment by heat loss or similar irreversible process. The update equation for the biomass of the microbe is thus given by Equation 12.2.

$$\Delta B = \theta C_{act} - \gamma \quad (12.2)$$

where C_{act} is the actual number of nutrient units consumed, which may be less than C_{max} if nutrients are scarce.

If an insufficient amount of a chemical nutrient is available to satisfy a microbe's consumption demand, the number of food units the microbe will try to consume is scaled down (possibly to zero) so that the demand can be satisfied. The ratio in which nutrients are consumed is always held fixed, so that the lack of a single nutrient can be a limiting factor to consumption. The limiting factor means that the actual amount of food consumed C_{act} is given by Equation 12.3 below.

$$C_{act} = \min(C_{max}, \frac{x_1}{\lambda_1}, \dots, \frac{x_M}{\lambda_M}) \quad (12.3)$$

where x_i is the amount of nutrient i present in the flask and λ_i is the proportion of consumption constituted by this nutrient. The amount of each nutrient removed from the environment is by a microbe is given by the vector E^- below, while the amount returned to the environment as excreta is given by the vector E^+ .

$$E^- = (\lambda_1 C_{act}, \dots, \lambda_M C_{act}) \quad (12.4)$$

$$E^+ = (\mu_1 C_{act}, \dots, \mu_M C_{act}) \quad (12.5)$$

For example, if the consumption and excretion ratios are $\lambda = (0.18, 0.5, 0.32)$ and $\mu = (0.3, 0.14, 0.56)$ respectively, and $\theta = 0.6$ the microbe will consume an amount of nutrients given by the vector $E^- = (1.8, 5, 3.2)$, gain 6 units of biomass, and excrete waste nutrient amounts of $E_+ = (1.2, 0.56, 2.24)$. If only 2 units of the second nutrient are available, the amount consumed from the environment will be $E^- = (0.72, 2, 1.28)$, the microbe will gain 2.4 units of biomass, and the amount excreted to the environment will be $E^+ = (0.48, 0.224, 0.896)$.

Microbe metabolism has a side-effect impact on the flask environment. For every unit of biomass created, the microbe will alter the level of each abiotic factor by the amount specified in its genotype, which may be positive or negative. The effect of the microbe's

metabolism on the abiotic environment is thus $(\alpha_1 C_{act}, \dots, \alpha_N C_{act})$. For example, if $N = 2$ and $\alpha = (0.2, -0.5)$, the microbe above that created 6 units of biomass would have an effect on the abiotic factors of $(1.2, -3)$.

In the experiments reported here parameter settings were $C_{max} = 10$, $\theta = 0.6$ and $\gamma = 1$, for all microbes.

12.3.2.3 Guilds and species

Given that microbes reproduce asexually, the normal usage of a species being a group of individuals capable of breeding with each other does not apply here, but for convenience we will refer to a microbe species as a group of individuals that share the same consumption and excretion ratios and have the same effect on the abiotic factors. Note that the normalisation process in the genotype-to-phenotype mapping means that a species can contain individuals with different genotypes.

Microbes can be grouped according to their functional effect on the environment. Through metabolism, every microbe has a (possibly negligible) effect on the amounts of different chemicals present in the environment. Ignoring the size of this effect and looking only at the sign of the effect (i.e., classifying effects as positive, negative or neutral), we can partition the microbial population into functional *guilds* for descriptive purposes.

The concept of an ecological guild was originally defined by Root as “...a group of species that exploit the same class of environmental resources in a similar way. This term groups together species, without regard to taxonomic position, that overlap significantly in their niche requirements” ([143], p.335). The concept has also been adopted by Volk, who defines biochemical guilds as groups of species linked by a similar effect on the environment, such as nitrogen fixers or oxygen producers [176], and also by Downing [55], who uses a similar definition.

Here we adapt the definition slightly so that a functional guild is made up of microbes that have a similar effect on the environment, e.g., a guild might contain all microbes who are net consumers of chemical A, net producers of chemical B, and increase the level of abiotic factor C. The guild classification allows some of the functional structure of the ecosystem to be elucidated. There are 3^{M+N} different functional guilds possible in the Flask model.

12.3.2.4 Reproduction

Microbes are initialised with a biomass randomly drawn from a uniform distribution on the range $[60, 110]$. If they are successful and increase their biomass to 120, they reproduce asexually by splitting. The child microbe gets a copy of the parent genotype and the

biomass of the parent is split equally between parent and child. The genotype of the child is subject to mutation, with a 0.05 probability of mutation at each allele. Mutation of an allele occurs by randomly selecting a new value from the range $[-1.00, 1.00]$.

12.3.2.5 Death

Unsuccessful microbes will not get enough food to grow and may reduce in biomass due to the cost of living. If the biomass of a microbe falls below 50 the microbe will die. In addition to this, a microbe may randomly selected to die with a probability of 0.002 per individual at each timestep. This is included as a loose implementation of death by natural causes and serves to promote selection and competition between microbes. When a microbe dies its remaining biomass is returned to the environment as chemical nutrients in equal amounts, i.e., $\frac{B}{M}$ units for each chemical.

12.3.3 Update scheme

The microbial population is updated sequentially in random order to avoid artefacts. At each timestep, each flask is updated in the following order:

1. Add chemical influx
2. While some microbes are not updated:
 - (a) Randomly select a microbe from the population
 - (b) Do microbe metabolism
 - i. Consume nutrients
 - ii. Update biomass
 - iii. Excrete waste
 - iv. Calculate effect on abiotic factors
 - (c) Test for microbe death (starvation or natural causes)
 - (d) Test for reproduction
3. Remove chemical outflux

12.3.4 Artificial selection experiments

There is no true concept of an ecosystem generation, but for convenience we will define a flask generation as a fixed period in which the flask ecology develops followed by a

selection event. Selection events occur periodically at fixed intervals, with the duration of a generation defined as the length of the between-selection interval. The length of this period was varied in different runs, and will be given where appropriate below.

For each evolutionary run a set of flux parameters was generated and held constant across all flasks to give identical flask conditions throughout each run. At initialisation a batch of flasks was instantiated with these parameters and then each was seeded with a randomly generated population of 100 microbes. At each selection event, a new batch of sterile flasks was seeded with 100 microbes randomly chosen (with replacement) from the highest fitness flask at the end of the previous generation, so that the fittest ecosystem provides the seed for all flasks for the next generation. Sampling was performed on the microbial population at the end of the growth period, i.e., from the ecosystem in the state it was in when its fitness was measured. A large number of evolutionary runs were performed with different flux parameters and the mean evolutionary response taken, in order to avoid unintended bias from favourable or unfavourable initialisation.

The ‘phenotypic’ ecosystem trait used for selection was based on the levels of the abiotic factors in the flask environment. An ideal vector \bar{Y} of target levels for each abiotic factor was pre-specified, with the deviation error of the state of a flask from the ideal vector constituting its fitness score. The error is given by Equation 12.6 below. Both the ideal vector and current state vector are normalised prior to calculation of error.

$$Error = \sqrt{\sum_{i=1}^N (\bar{y}_i - y_i)^2} \quad (12.6)$$

This fitness metric was chosen for similarity with Swenson et al’s selection for pH level in their pond-water ecosystems [163], and also because it does not affect the relative fitness of individual microbes within a flask ecosystem. For a single flask the state of the environment is a free parameter with respect to within-flask evolution, meaning that any non-random variation in environmental composition can be attributed to higher-level selection between flask ecosystems.

12.4 Results

12.4.1 Ecosystem dynamics

First of all we briefly describe the ecological dynamics that occur within a flask ecosystem. Figure 12.1 shows the guild structure of a typical ecosystem in a flask where $M = N = 2$. This is a very simple scenario, chosen to aid understanding.

The shaded striations in Figure 12.1 represent the number of individuals belonging to

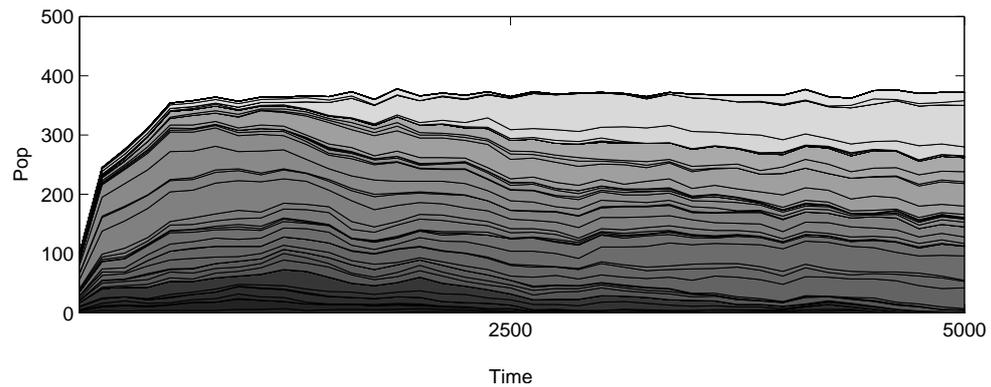


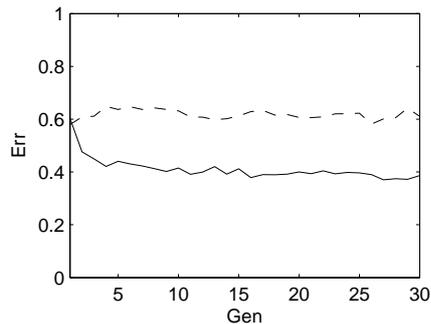
Figure 12.1: Ecological dynamics of a single flask ecosystem. Shaded striations denote size of different functional guilds, while the stacked height of the plot denotes the total population size. The population grows quickly to a stable size. The distribution of individuals between different functional guilds is initially quite uniform, but by the end of the simulation run several large guilds have developed, indicating the formation of a food web.

different guilds. These quantities are stacked, so that the total height of the plot gives the size of the whole microbial population. The flask is seeded at $t = 0$ with 100 microbes with randomly generated genotypes. These initially reproduce rapidly until the population reaches its carrying capacity, so that the size of all the guilds increases steadily for the first few hundred timesteps before levelling off.

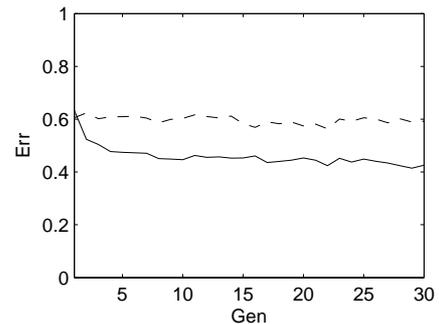
As time passes, reproduction and mutation cause the emergence of new species (sometimes reflected in the appearance of new guilds). Ecological interaction between individuals means that certain species may do well and dominate at different times. The guild structure reflects this, with new guilds emerging and growing, or shrinking and disappearing. The distribution of individuals between guilds is much less even at the end of the run, with a few large guilds and many small guilds indicating the formation of a food web, as compared to the more uniform early distribution resulting from the random initialisation.

12.4.2 Response to selection

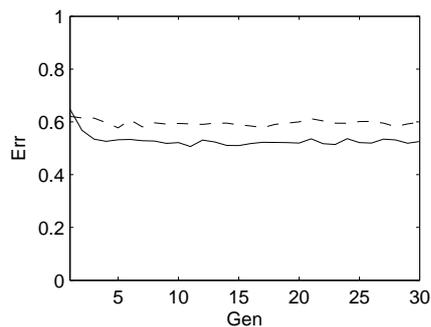
Figure 12.2 shows the response of the flask ecosystems to artificial selection. The plots show fitness increasing (i.e., error reducing) over generational time when ecosystems are selected for minimal error in the levels of abiotic environmental factors. The data are generated from runs where $M = N = 3$ with a normalised ideal vector for abiotic factors of $\bar{Y} = (0.2, 0.3, 0.5)$; similar results to those presented here are obtained when these parameters are varied but are omitted here for clarity. The duration of each ecosystem generation in timesteps was varied, with runs performed for all values in the set



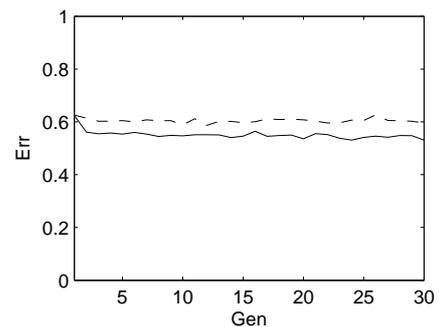
(a) 1000 timesteps (60 directed/ 55 random)



(b) 2000 timesteps (54 directed/ 52 random)



(c) 5000 timesteps (59 directed/ 51 random)



(d) 10000 timesteps (58 directed/ 57 random)

Figure 12.2: Mean fitness over generational time for ecosystem generations of different durations. Solid lines are results for directed selection, dashed lines are results for random selection. Numbers in brackets denote number of runs performed with directed and with random selection; plot shows mean values over all runs. Divergence of the solid line from the dashed line indicates a response to selection.

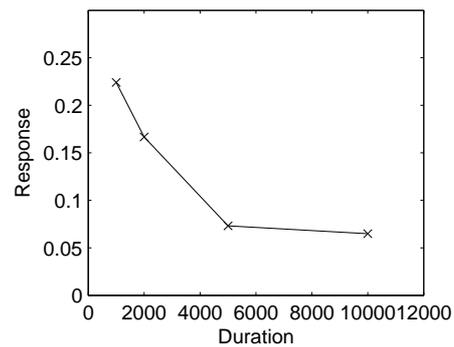


Figure 12.3: Response to selection (difference between directed-selection fitness and random-selection fitness) after 30 generations for ecosystem generations of different durations. Response is reduced as the length of the period between selection events is increased.

{1000, 2000, 5000, 10000}. For comparison, a microbe will reproduce every 10 timesteps if there are no food constraints, so all of these generation durations allow at least 100 microbial reproductive cycles. More than 50 separate runs were performed with each different generation length, to give a reasonable data sample.

In each plot the response to selection for minimal environmental error is plotted against a control group of runs where parent ecosystems were chosen at random from the population of flasks. Random selection in the control runs allows us to see how the system would behave in the absence of selection, and gives a baseline for comparison with the directed selection runs. The mean fitness of all flasks in each generation is plotted to show the trend of the whole population of flasks, rather than focusing on one (possibly anomalous) individual flask. It is clearly apparent that in all cases there is a significant response to directed selection, with error values falling steadily over generational time.

12.4.3 Time-dependence of the response to selection

It is apparent from Figure 12.2 that the response to selection is reduced as the length of the generations is increased. This can be seen more clearly in Figure 12.3, which shows the difference in mean fitness after 30 generations between the directed selection runs and the random-selection control runs, plotted against the duration of each generation. The reason for this inverse relationship between the response to selection and the duration of each generation can be explained by consideration of the evolutionary and ecological dynamics within each individual flask.

The processes of ecological competition and evolutionary change are continuous, meaning that the ecosystems never reach equilibrium. Experiments (not shown here)

have shown that the flask ecosystems do not reach an attractor (fixed-point, limit cycle or chaotic) even when left to run for up to 5×10^5 timesteps. This does not mean that they are without structure, as the functional guild distribution may remain broadly similar for very long periods, but that the underlying genetic change in the population is unceasing.

This continual change means that the microbial population moves steadily away from its initial composition. So if a flask is seeded with a particular set of microbes, the similarity of its species distribution to its initial state steadily decreases. In the current scenario, there is ecological pressure for the maintenance of a stable food web, but there is no constraint on the effect of the microbe population on the environment. Thus genetic drift can erode the species distribution that was favourable to reducing the environmental error and increasing fitness. Since the similarity of the sample population taken at the end of an ecosystem generation to the initial seed population falls as the duration of the generation increases, the erosive effect is greater with longer generations. This means that the likelihood of a fit ecosystem being maintained in the face of genetic drift long enough to be passed on to the next generation decreases as the generation length increases; heritability is reduced by longer generations.

12.5 Discussion

We have presented an evolutionary simulation model in which microbial ecosystems display a response to artificial selection on an ecosystem-level trait. This response is decreased when the duration of an ecosystem generation is increased, due to genetic drift in the alleles governing the effects of microbes on the flask environment.

The collective effects of these environment-altering traits of individual microbes determine the levels of abiotic environmental factors and form the ecosystem trait that is selected for. Genetic drift in these alleles occurs because although ecologies containing good sets of environment-altering genes are selected for at the ecosystem level, the environment-altering alleles have no effect on the relative fitness of microbes within a flask and so vary randomly during within-flask individual-level selection.

The longer the ecosystem generation, the greater the amount of genetic drift that occurs. Since offspring ecosystems are seeded with samples taken from parent ecosystems at the end of each generation, longer generations reduce heritability and reduce the adaptive response at the ecosystem level.

Our findings agree with those of Swenson et al (2000b) and also suggest a new experiment that could be performed in the laboratory. The time-dependent nature of the ecosystem-level response to selection could be explored using microbial ecosystems with a similar experimental method to that described in [162, 163].

With reference to the Flask model, these experiments demonstrate the possibility of heritability between parent and offspring flask communities. A sustained response to selection was observed and this implies a non-trivial level of heritability. The dependence of the response to selection on the length of the period between selection events is an interesting feature of the system, but should not cause any problems in the Flask scenario; in the Flask model selection is an implicit function of migration, colonisation and invasion, all of which are continuous processes.

12.6 Conclusion

In this chapter we set out to demonstrate heritability in the Flask model. In a simple implementation we have shown that a response to selection at the level of a flask ecosystem can reliably be observed. A sustained response to selection implies a measure of heritability, so our aim has been reached.

The next step in the development of the Flask model is to demonstrate that when feedback from the environment to microbial growth is present, a correlation exists between the composition of the environment and the size of the microbial population in a flask. Once this has been done, experiments can be performed on spatially structured arrays of interconnected flasks, in order to test whether or not environmental regulation evolves by multi-level selection.

Chapter 13

Discussion of Part IV

13.1 Overview

This chapter briefly summarises the main findings from Part IV, and tries to analyse the main contribution of this work for Gaia theory and biology in general. The methodology used is critically reviewed, and the most promising directions for future work are outlined.

13.2 Summary of results

13.2.1 Chapter 9: Background

Part IV of this thesis is an investigation into the possibility of environmental regulation by the evolving biota. The idea of environmental regulation is most commonly explored under the mantle of Gaia theory, and the majority of the work presented in Part IV is in this area. However, it is hoped that the generality of the models and theoretical arguments discussed in Part IV is such the results achieved are not specific to the Earth system, but will also allow useful insights to be gained concerning environmental regulation at a variety of other scales.

The basis for Gaia theory is the observation that the feedback between organisms and their environment works in both directions; the environment sets the conditions in which organisms adapt and evolve, but organisms also have an impact on the physical environment. The Gaia theory claims not only that this double feedback occurs, but that

it leads to regulation of the environment by the biota, in conditions suitable for life.

The Gaia hypothesis was originally put forward by Lovelock and Margulis [124] and stirred up significant debate in the literature, much of which centred around the compatibility of Gaia theory with Darwinian evolution. The big question raised by evolutionary theory concerned how Gaian regulatory feedback loops could be created by selfish natural selection.

The case against Gaia was begun by criticisms of teleology and vulnerability of a global regulatory system to ‘cheaters’ [40,52], with the modern view well-put by Kirchner [98]. Kirchner accepts the emergence of global feedback loops but feels that there is no reason why such feedbacks should be regulatory; why shouldn’t anti-Gaian positive feedbacks be formed? Kirchner also points out that the existence of a good fit between organisms and their environment is easily explained by adaptation of organisms to their environment, and that there is no need to posit any additional mechanisms.

Lenton [105] exemplifies the case in favour of Gaia, which is in part evidential (regulatory feedbacks *must* have evolved because the chances of life persisting on Earth as long as it has without regulation are infinitesimally small) and in part theoretical (Daisyworld shows that global regulation is not incompatible with selfish individual selection, and cheaters would not disrupt regulation since it occurs as a by-product of organismic activity).

The latter section of Chapter 9 describes the Daisyworld model [183], which was first developed to refute claims of teleology, but has since been the subject to many expansions. The Guild model [55] is also described, since it represents the only significant attempt in the literature to move on from Daisyworld.

Finally, Chapter 9 put forward a list of research questions to be addressed in the remainder of Part IV. It was noted that in both Daisyworld and the Guild model, organisms can create a local buffer against the environment and selfish adaptation always contributes to global regulation; there is no possibility of locally advantageous mutations that would destroy regulation. One aim for Part IV was therefore to develop a good understanding of the conditions under which environmental regulation evolves in Daisyworld, and to examine the consequences of relaxing the core assumptions of the model. A further aim was to develop a new model based on different assumptions to Daisyworld, in order to investigate the evolution of biotic regulation in a different scenario.

13.2.2 Chapter 10: Evolution in Daisyworld

Chapter 10 presents a simplified version of Daisyworld developed with the purpose of studying the effects of Darwinian evolution in the Daisyworld scheme. The physics of the

model are simplified (following Harvey [80]) in order to aid understanding of its dynamics, while its implementation as an individual-based model (rather than the population-based form of the original Daisyworld) allows easy inclusion of Darwinian evolutionary adaptation. This model was used to gain an understanding of the conditions under which environmental regulation will evolve, and also to test the effects of relaxing the core Daisyworld assumptions.

The simplified Daisyworld model consists of a toroidal cellular automata where each location (or patch) can be bare or occupied by a daisy species. Bare patches have a neutral albedo, while daisy species occupying a patch can raise or lower patch albedo by their pigmentation. The albedo of a patch determines the amount of heat it absorbs from solar insolation, and since daisies grow best at a particular temperature, albedo affects growth. The preferred growth temperature and albedo of a daisy species are subject to mutation when daisies reproduce.

First of all it was demonstrated that the simplified model could replicate the behaviour of the original Daisyworld model, where competition between the black and white daisy species led to regulation of the global temperature in the face of steadily increasing solar insolation. Known Daisyworld results involving evolutionary adaptation were also replicated. In a series of experiments it was found that allowing daisy albedo alone to mutate freely between the levels for black and for white daisies led to improved regulation of global temperature. However, allowing the preferred growth temperature to mutate freely led to the collapse of regulation, as daisies simply evolved to prefer the current ambient temperature.

These experimental results suggested that regulation would evolve where there was both a need to regulate (i.e., a constrained growth function) and an ability to regulate (i.e., variation in albedo). To test this hypothesis further experiments were performed in which daisy albedo could mutate freely, but where mutation of preferred growth temperature was restricted. The results achieved confirmed the hypothesis. When only two growth temperatures were possible, regulation occurred first around the lower temperature, then around the higher temperature, as solar insolation increased. When preferred growth temperature was allowed to mutate freely, but at a much slower rate than mutation of albedo, regulatory epochs occurred around different preferred growth temperatures.

Having shown the importance of constraints on evolutionary adaptation for regulation to occur, the next set of experiments looked at the effects of relaxing two key assumptions of the Daisyworld model, which were the assumptions that daisies could create a local patch temperature different from the global temperature, and that locally beneficial adaptations always contributed to global regulation. Both assumptions were found to be vital for regulation to occur, with a failure of regulation if either assumption was relaxed.

Chapter 10 concluded with the observation that all of the criteria for Daisyworld regulation could be met in the natural world, but that the creation of a local buffer against the environment was likely to be problematic. It was also noted that the simplified Daisyworld model that had been developed was more general than the original Daisyworld model, meaning that the results obtained might be applicable in a variety of different scenarios and at a variety of scales.

13.2.3 Chapter 11: The Flask model

Chapter 11 describes a new and original model, the Flask model, which is intended to look at the evolution of environmental regulation by the biota in a situation where some of the key assumptions of the Daisyworld models are relaxed. Specifically, the Flask model will be used to study whether or not biotic regulation can evolve when it is possible for selfish individual adaptations to work against global regulation, and where individual organisms cannot create a local buffer against the environment.

The Flask model consists of an array of flasks containing a liquid medium and supplied with a constant flow of nutrients. Each flask contains a population of microbes, which can grow, reproduce and mutate, so that stable food webs may be formed by ecological and evolutionary interaction. Microbes have an effect on their abiotic environment, and it is assumed that the abiotic environment has an effect on growth, so that a two-way feedback exists between the microbial community and the abiotic environment in which the community determines the composition of the abiotic environment and the abiotic environment influences the rate of microbial growth.

There is limited migration between flasks, allowing the possibility of colonisation or invasion of a neighbouring flask by a community. We would expect more populous communities to be better colonisers, and since the effect of a microbial community on its abiotic environment will be key to its growth and population size, it is possible that higher-level selection effects will lead to the evolution of communities that regulate their environment around the optimal conditions for growth.

13.2.4 Chapter 12: Response to ecosystem selection in the Flask model

Chapter 12 presented a model designed to investigate the possibility that heritability could exist between a 'parent' community and its colonist 'offspring'. It was argued that if a response to selection at the level of a flask community could be demonstrated, then this would constitute proof of some degree of heritability. The aim of the chapter was to implement a simple version of the Flask model and show that a response to selection could be observed.

Simulated flasks were prepared as in the Flask scheme described in Chapter 11, but with the simplifying assumptions that all flasks were isolated from each other, and that there was no feedback from the abiotic environment to microbial growth. The composition of the abiotic environment of a flask was used as a 'phenotypic' trait at the level of the flask community, and a target composition vector was defined; the error between the actual composition of a flask and this ideal was the fitness metric used for artificial selection experiments on a population of flask communities.

It was found that a clear response to selection was observed, with the mean environmental error in the population of flasks steadily reducing over generational time. A clear deviation from the performance of a random-selection control group was observed. It was also found that the response to selection was inversely proportional to the length of time for which each flask community was allowed to develop between selection events.

The conclusion to Chapter 12 argued that since a response to selection at the level of the flask community had been demonstrated, heritability must exist between parent and offspring communities. The time dependent nature of this result was not thought to be a problem for the Flask model, since in the Flask model migration and colonisation would be a continuous process.

13.3 Implications for Gaia theory

The simplified Daisyworld model presented in Chapter 10 is a useful addition to the Gaia literature because it simplifies the physics of the Daisyworld and allows the key features of the model to be seen more clearly; a semi-realistic implementation of heat radiation adds little to our understanding of the compatibility of individual selection and global regulation, whereas a simple model of the mutation of preferred growth temperature can teach us a lot. Using our simplified model we have gained an understanding of the circumstances under which we can expect to see regulation in the Daisyworld model. The simplicity of our model makes it more general than the original Daisyworld, so our findings should hold in more 'realistic' Daisyworld models and it is also possible that they can be applied to other systems that share a similar form.

We found that regulation in Daisyworld will occur when there is heritable variation in daisy albedo and constraints on the adaptation of the preferred growth temperature. Variable albedo allows the possibility of regulation, while constraints on adaptation of growth temperature creates a need for regulation. This finding might guide field research aimed at finding examples of real-world regulation that operates along similar principles to Daisyworld. There are many constraints on evolutionary adaptation in nature, arising from physical/chemical laws and from evolutionary canalisation. Real world ecologies

where such constraints on adaptation are identified would be good places to look for regulation.

Our examination of the impact of constraints on adaptation of preferred growth temperature involved two results not previously reported in the Daisyworld literature. The first of these is that when there are two possible growth functions with overlapping ranges but distinct maxima, competitive exclusion effects are observed whereby biotic regulation of the environment around one of the growth maxima precludes organisms with metabolisms based on the other growth function from gaining a foothold in the environment. The second new result is that when adaptation of preferred growth temperature occurs without constraint, but very slowly in comparison to adaptation of albedo, regulatory epochs are observed in which the biota regulate around a particular dominant growth function for a period, before changing external conditions make that growth function unviable. In the general case this means that slow adaptation of a phenotypic trait affected by the environment coupled with fast adaptation of a trait affecting the environment can lead to short-term adaptation *of* the environment interspersed with infrequent adaptation *to* the environment. Both of these results are novel in the Daisyworld literature, and suggest possible new avenues for Gaian research.

Other significant new results concerning Daisyworld concern the effect of relaxing its key assumptions. We found that if organisms were not able to create a local buffer against the environment in which to improve their growing conditions, regulation broke down. There was no individual selective advantage to adaptations that improved the environment, so although occasions could arise when serendipitous circumstances led to the right kind of daisy improving the environment and the population expanding rapidly, there was no selective pressure to maintain this improvement and it was quickly eroded by genetic drift. This finding shows that there is no way that individual-level selection can lead to sustained global regulation unless there is some individual-level competitive advantage to be gained by contributing to it, which lends support to criticism of Gaia hypotheses based on individuals *altruistically* contributing to regulation.

Also, if the effect of daisy albedo on the environment had no effect on growth so that mutations of albedo were selectively neutral, regulation also broke down. This shows that if the direct link between selfish adaptation and global regulation is relaxed in the Daisyworld model, then regulation fails. This finding echoes Kirchner's comment that there is no reason to expect biological feedbacks to the environment to be inherently homeostatic.

It therefore seems that regulation in Daisyworld rests on the twin assumptions of a local environmental buffer and of selfish adaptation contributing to the global good. These are reasonable assumptions to make in many real world situations, but they are by no

means universally true. Given the importance attached to the Daisyworld model in the Gaia literature (notwithstanding the fact that the validity of Gaia theory does not rest on the validity of the Daisyworld model) the importance of these assumptions for the occurrence of regulation in Daisyworld should be given greater recognition than is currently the case.

However, as mentioned, the Gaia theory does not rest on Daisyworld, and in response both to our findings concerning Daisyworld and to Volk's call for new Gaia models, we have in this thesis begun to develop a new model, the Flask model. It is hard to judge the implications of a model which has yet to be constructed, but the fact that the Flask model is built on different assumptions than the Daisyworld model means that it can provide another test case for the evolution of environmental regulation and move the Gaia debate forward.

The Flask model may appear to be swapping one controversial theory for another in calling on multi-level selection to demonstrate the validity of the Gaia theory, and to some extent this is true. Group selectionist explanations of global regulation have previously been refuted, but we feel that recent support for the validity of multi-level selection in certain circumstances suggests that such explanations, carefully constructed, are worthy of further investigation.

13.4 Implications for biology

The stated aim of the investigation in Part IV of this thesis was to study the evolution of environmental regulation by the biota, and although the large majority of the material presented above is aimed at Gaia theory, this aim of generality still applies. It seems likely that biotic environmental regulation can occur at a variety of spatial and temporal scales, and it would be beneficial if the results obtained above could be applied to ecosystems other than the Earth system.

The evolutionary Daisyworld model presented in Chapter 10 is made very general by its simplification, as mentioned before. Thus it might be a good starting point for understanding environmental regulation in systems other than Gaia; if the assumptions it is based on can be shown to hold in a candidate regulatory system, and the criteria identified for regulation to occur are met, then useful insights might be gained by comparison of the dynamics of the target system with those of the simplified Daisyworld model.

The Flask model can also be applied at a variety of scales. Its conceptual design stems from microbial activity in a liquid environment, but it might be adapted to model other situations without much alteration to its structure; it is a thought experiment more than a model of a particular ecosystem.

In addition, the Flask model can be used to address questions other than those concerning environmental regulation. Multi-level selection is a supporting assumption of the Flask model as described in Chapter 11, but it is in itself an issue of great theoretical importance and one for which the basic formulation of the Flask model is well suited for investigation. Experiments might include a more rigorous and thorough exploration of the response to artificial selection presented in Chapter 12, or perhaps a study of whether or not the spatial structure of the Flask model provides necessary and sufficient conditions for multi-level selection to occur in a more natural form.

Finally, because the Flask model is loosely based on microbial communities inhabiting flasks, it also offers the intriguing possibility of making predictions that might be tested with laboratory experiments using real microbes. The similarity of the results from simulation presented in Chapter 12 to those of Swenson et al working with pond-water and soil communities [162, 163] suggest that this might be a useful line of study. For example, it would be interesting to see if the time-dependent nature of the response to ecosystem selection in simulation is a property shared by real microbial communities. The possibility of real world testing of hypotheses generated from the Flask model is one area where it may offer an advantage over the less directly testable results generated in Daisyworld.

13.5 Review and future work

The main area where more work is required is in developing the Flask model. The lack of time and space to provide a more concrete development of the Flask model in this thesis is a major regret. However, despite the cartoon nature of the Flask model as described, and while accepting fully that the model will inevitably be subject to myriad refinements and corrections, we feel that its inclusion here is justified. The inclusion of the Flask model adds balance to Part IV of the thesis. Chapters 9 and 10 point to a direction in which Gaian modelling might usefully move forward, while the limited presentation of the Flask model in Chapters 11 and 12 gives a taste of what such a journey might bring.

Future work will focus on developing the Flask model. The assumptions and sub-hypotheses described in Section 11.6 in Chapter 11 will be examined in turn. The demonstration of heritability and response to artificial ecosystem selection that was presented in Chapter 12 will be extended and better analysed. Further modelling work will be required to demonstrate a correlation between the population size of a community and its effect on its abiotic environment, and to show that the spatial structure of the Flask model can give rise to multi-level selection dynamics. A literature search will be necessary to compile further evidence for the existence of universally preferred growth conditions, and this

might be complemented with modelling to investigate the implications of multiple coexisting sets of preferred environmental conditions. Only when these components of the Flask model are studied and understood can any principled testing of the overarching hypothesis concerning the evolution of environmental regulation by the biota be performed.

13.6 Conclusions

Part IV of this thesis has considered the evolution of environmental regulation by the biota. A simplified version of the Daisyworld model was used to gain a good understanding of when regulation can be expected to evolve in the Daisyworld scenario. The effect of relaxing the main assumptions of Daisyworld was also studied, and it was found that regulation collapses when some of these modelling assumptions are removed. The Flask model was suggested as a system that relaxes the assumptions of Daisyworld but in which environmental regulation might be expected to evolve by multi-level selection. Finally, Part IV presented the first steps towards constructing the Flask model, and demonstrated that at least one of its component hypotheses holds.

The contributions of this body of work for Gaia theory and for biology in general have been outlined above, along with known weaknesses of the material presented, and a plan for future work. This concludes the study of environmental regulation by the biota in this thesis. In Part V we will examine biotic regulation of the environment in the context of homeostatic adaptive networks.

Part V

Discussion

Chapter 14

Towards a general theory of homeostatic adaptive systems

14.1 Overview

The aim of Part V is to try and draw the various threads of research in the thesis together and reach some coherent conclusions concerning homeostatic adaptive networks. In Part I we introduced the concept of homeostasis as an organising principle, and in Part II we described an example system, the Homeostat. Part III was a study of homeostatic plasticity in neural networks, while Part IV examined environmental regulation by the biota. In Part V we return to the central theme of homeostatic adaptive networks and analyse each of the systems presented in this thesis in this context. The aim is to see how well the Homeostat, neural homeostasis, and biotic environmental regulation fit into our definitions of homeostasis and homeostatic adaptation, and to try to identify properties common to all homeostatic adaptive systems.

The next section gives a brief recapitulation of the definitions of different kinds of homeostatic adaptation that we originally gave in Chapter 1. These definitions are then applied to our case study systems in subsequent sections. This is followed by some discussion of commonalities between the case study systems and general conclusions concerning homeostatic adaptation.

14.2 Definitions

In Chapter 1 we defined a homeostatic system as one where the essential variables were maintained in a target range by the action of a regulatory mechanism, where they would otherwise have moved outside the target range due to external forcing. In Figure 1.4 we gave a schematic of this organisation.

We then defined homeostatic adaptation as changes in system behaviour that maintain homeostasis in essential variables against some perturbation. Elastic homeostatic perturbation is where these changes are not persistent, i.e., the system returns to its original state after perturbation. Plastic homeostatic adaptation refers to the case where changes are persistent, i.e., where the system remains in a changed state after the perturbation is removed.

We then noted that it is possible for multiple homeostatic feedback loops to operate simultaneously in the same system (Figure 1.5) and that second order homeostasis is possible where homeostatic regulation acts to change the parameters of a first order homeostatic regulator (Figure 1.6).

14.3 Case study 1: The Homeostat

The original Homeostat was an electromechanical device consisting of 4 units interconnected by circuits with variable parameters. Each unit had a magnet mounted on its upper surface. The angular deviation of each magnet was determined by an electrical field that varied as a function of the input current received by the unit, modified by the parameters of a set of potentiometers and commutators. If the angular deviation of a magnet passed outside the range $[-45^\circ, 45^\circ]$ a uniselector randomly chose new values for the potentiometer and commutator parameters, so that the system was only stable when all magnets were inside the target range.

External perturbations to the Homeostat come in the form of externally induced displacements of one or several of the magnets on the four units. The target variables of the system are the angular deviations of the magnets. Some perturbations to a magnet will act on the underlying circuitry so that the system falls to a new attractor with all magnets in bounds, and homeostasis is maintained. Some other perturbations will lead to system instability with some magnets out of bounds, i.e., a loss of homeostasis. The latter case will activate the uniselector to choose new values for the potentiometer/commutator settings of the affected unit, which may alter the properties of the circuit so that all magnets are returned to their target ranges and homeostasis is restored.

The simulated Homeostat presented in Chapter 2 behaves similarly to the original

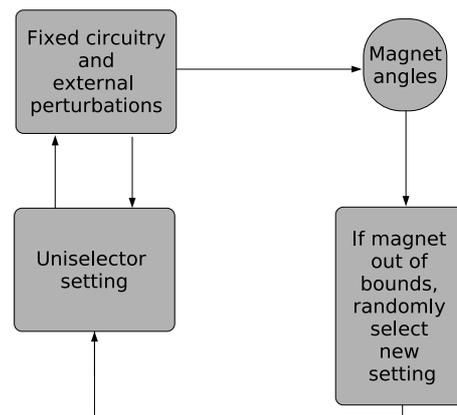


Figure 14.1: The original Homeostat is a first order plastic homeostatic adaptive system. Fixed circuitry and external perturbations combine with the current uniselector settings to determine the angular deviations of the magnet mounted on each unit. If a magnet is outside the target range, a new uniselector setting for that unit is chosen at random.

Homeostat. Here the target variables are the activation values for each unit, which are determined by the network structure, connection strengths, and unit transfer function parameters. Perturbations come in the form of external input signals to each unit. Some perturbations can be accommodated by the simulated Homeostat without any loss of homeostasis, but others cause activation of one or more units to go out of bounds, triggering the random generation of new connection strengths and transfer function parameters for non-homeostatic units.

Applying our definitions to both the original and simulated Homeostats, we see that they are both first order homeostatic adaptive systems. They both display plastic homeostatic adaptation, and there are different homeostatic feedback loops regulating each unit so that the whole 4-unit system has multiple interacting homeostatic subsystems. The simulated Homeostat also has two different regulatory mechanisms active on the same target variable, which are the two mechanisms for reassigning connection strengths and transfer function parameters. We can therefore draw the diagrams of causal effect shown in Figures 14.1 and 14.2.

It is interesting to note that the decision of where to draw the boundary between the reacting part of the system and the environment seems very arbitrary. External perturbations are clearly part of the environment, and the current uniselector setting is clearly part of the reacting part, but other parts are harder to ascribe to one or the other. For example, the fixed circuitry of the Homeostat part of the environment or a part of the reacting part? Here we will adopt the principle that the reacting part consists of those parts of the system that determine the type of behaviour the system produces, and that the environment consists of everything between the reacting part and the essential variables. This refinement

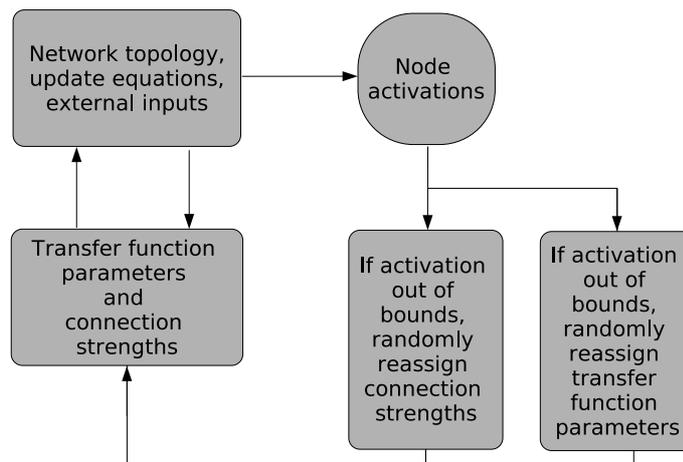


Figure 14.2: The simulated Homeostat is a first order plastic homeostatic adaptive system with two homeostatic feedbacks. Network topology, update equations and external inputs combine to determine the activation of each node. If activation of a node is out of bounds, adaptive mechanisms randomly reassign transfer function parameters and afferent connection strengths for that node.

of the definition treats the environment as a black box which converts the current state of the reacting part into a set of values for the essential variables. Hence the reacting part in the original Homeostat is the uniselector, and the environment consists of all the fixed circuitry and any external perturbations. In the simulated Homeostat, the reacting part is the current set of connection strengths and transfer function parameters, while the environment is the fixed set of network update algorithms and any external inputs.

14.4 Case study 2: Homeostatic plasticity in neural networks

We applied homeostatic plasticity to continuous-time recurrent neural networks (CTRNNs) in the form of two different mechanisms that were activated when neural firing rate was too high or too low. These mechanisms were local (in that they acted on the parameters of the non-homeostatic node) and directed (in that they always changed parameters in the correct way to move firing rate towards its target range). The synaptic scaling mechanism caused a linear scaling of afferent connection weights to a neuron so that input caused more or less excitation, while the adaptive bias mechanism changed the bias term so that the neuron was inherently more or less excitable. Both synaptic scaling and adaptive bias were effective in acting on node parameters so that extreme firing rates were counteracted and firing rate homeostasis maintained.

The target variables in this system are the firing rates of the nodes, with a target range

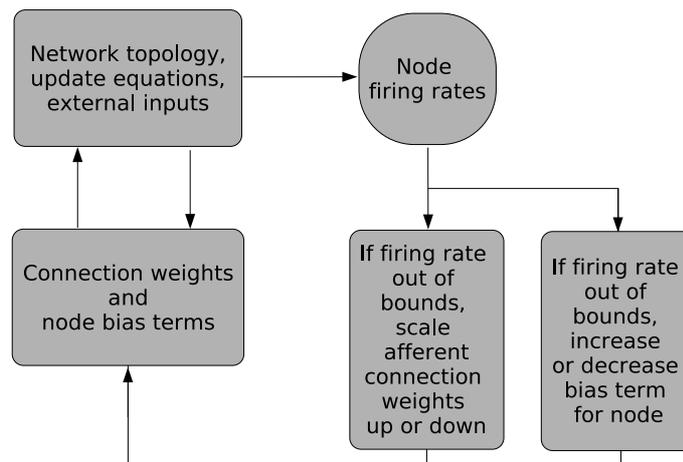


Figure 14.3: Homeostatic feedback loops in homeostatic plastic CTRNNs. Homeostatic plastic CTRNNs are first order plastic homeostatic adaptive systems with two regulatory mechanisms: synaptic scaling and an adaptive bias term. The current set of connection weights and bias terms combines with the network topology and update equations to determine the firing rates of the nodes in the network. If the firing rate of a node is too high, afferent connection weights are scaled down and the bias is increased, reducing the firing response. If firing rate is too low, weights are scaled up and bias reduced, making the node more likely to fire.

(e.g. $[0.3, 0.7]$) that is a subset of the full range $[0, 1]$ permitted by the sigmoidal form of the CTRNN transfer function. Perturbations to a CTRNN take the form of external input signals applied to the nodes. After initialisation, there is a transient period in which homeostatic plasticity acts on network weights and bias terms so that the activity of each neuron is brought inside the target range. After this transient, most input signals will not cause activity to go out of bounds and there is no parameter change during normal operation. However, certain input signals will lead to patterns of activation where neural activity goes out of bounds, at which point the homeostatic plastic mechanisms will act so that activity is returned to the target range.

From our definitions, we can state that a homeostatic plastic CTRNN is a first order homeostatic adaptive system and that it displays plastic adaptation. There are multiple homeostatic subsystems regulating the activity of each of the nodes. Synaptic scaling and adaptive bias are two alternative types of regulatory mechanism acting on the same variables with the same target range. A diagram of causal effect for a network with both mechanisms active is shown in Figure 14.3. Again the division between environment and reacting part seems somewhat arbitrary, but following the principle that we applied to the Homeostat we assign the current set of network weights and bias terms to be the reacting part and the set of network update equations and external inputs to be the environment.

The complexity of biological neuronal networks compared to CTRNNs makes it hard to draw any conclusions concerning the nature of homeostatic plasticity in biological

nervous systems from our results in this thesis, but a similar scheme to that shown in Figure 14.3 might apply.

14.5 Case study 3: Homeostatic regulation of the environment by the biota

Since the different models (Daisyworld and Flask) that are described in this thesis are different in nature from each other and from the Earth system that inspired them, each will be treated separately.

14.5.1 Daisyworld

Several versions of the Daisyworld model were examined, which were variations on a basic scheme where daisies compete for living space on a toroidal landscape. Each patch contains a single daisy species or bare earth, and its resulting albedo affects the amount of solar heat radiation it reflects and hence its local temperature. Daisies may colonise bare neighbouring patches, with the possibility of mutation of daisy albedo and/or preferred growth temperature during this reproductive process. Ecological competition and evolution result in regulation of global temperature.

Since some versions of the Daisyworld did not display any regulatory properties, they are not candidate homeostatic adaptive systems. Here we will look only at the Daisyworld scenarios where homeostatic regulation of global temperature was observed. The target variable of the Daisyworld system is always global temperature, with perturbation taking the form of a steadily increasing level of solar luminosity.

14.5.1.1 Two-species Daisyworld

The behaviour of the Daisyworld system in this scenario is shown in Figure 10.2 and described in Chapter 10. When the Daisyworld is configured with pure black and pure white daisy species only, with no mutation, regulation occurs because of a shifting balance between the sizes of the black and white daisy populations. This balance arises as a result of competition between the daisy species for living space, and interaction between daisies and their environment.

The two-species Daisyworld is a first order homeostatic adaptive system, but in this case the adaptation is elastic. Regulation is achieved by a balance between two species that are always present in the world; it is the relative sizes of their populations that determine temperature. Population sizes tend asymptotically to fixed values for any given

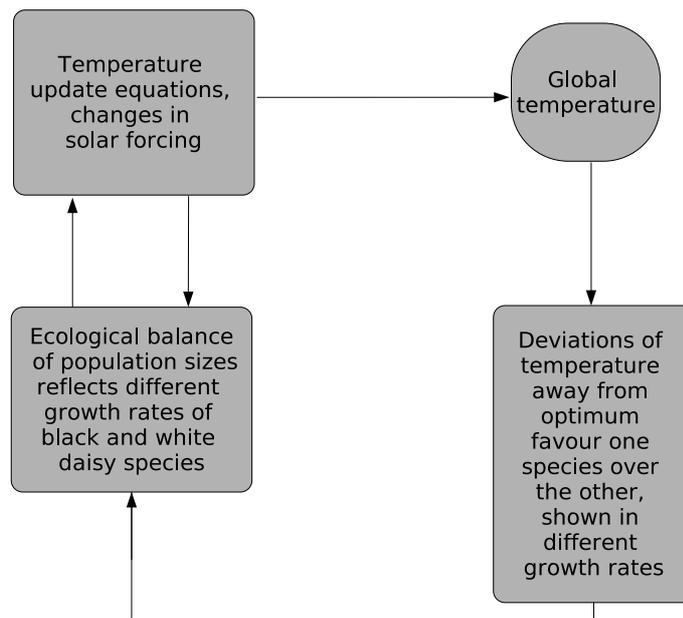


Figure 14.4: The two-species Daisyworld is a first order elastic homeostatic adaptive system. If global temperature is higher than the optimum level for growth, white daisies have a higher growth rate than black daisies; vice versa when global temperature is lower than the optimum. Ecological dynamics determine the population sizes of the black and white daisy species, which in turn combine with solar forcing to determine the global temperature. Differences in growth rate affect the ecological dynamics, leading to temperature regulation.

level of external forcing, so in the limit the state of the system is determined entirely by current conditions. Therefore any changes in population sizes caused by a change in external forcing will persist only as long as that level of forcing continues, and changes are reversible. The system displays elastic homeostatic adaptation since changes in response to a change in forcing do not persist.

We can draw the diagram of causal effects for the two-species Daisyworld that is shown in Figure 14.4. A deviation in temperature away from the optimal level for daisy growth creates an advantage for one species over the other that is expressed as a higher growth rate for the species most suited to the current temperature. Daisy ecology alters the balance of population sizes, which in turn moves the global temperature towards the optimal level.

14.5.1.2 Evolutionary Daisyworld

Short-term regulation occurs in the evolutionary Daisyworld scenario (see Figures 10.3, 10.6, 10.7 and 10.10, and discussion in Chapter 10) in a similar way to the two-species model. Ecological interaction and competition between different daisy species results in a balance of population sizes that gives a global mean albedo sufficient to move global

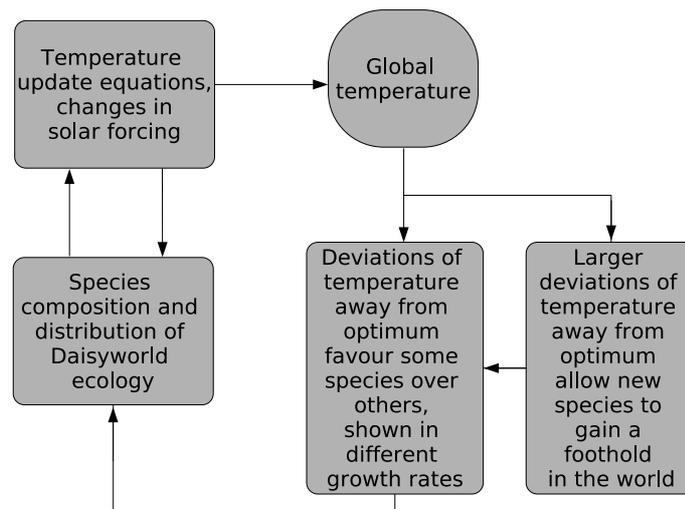


Figure 14.5: The evolutionary Daisyworld is a second order homeostatic adaptive system. Ecological interactions form a first order elastic homeostatic adaptive mechanism which can regulate small perturbations to global temperature. Larger perturbations create an opportunity for new species to enter the world and contribute to regulation. Evolutionary adaptation thus forms a second order plastic homeostatic adaptive mechanism that acts on the parameters of the first order ecological mechanism.

temperature to the optimum for growth.

In the two-species model there are only pure black and pure white daisy species, meaning that while it is possible to regulate the global temperature, the local growing conditions for each daisy species is often non-optimal and the combined population size rarely reaches carrying capacity (see Figure 10.2 in Chapter 10). In the evolutionary Daisyworld, mutation continually adds new daisy species to the daisy ecology, meaning that there is usually some variant that can create optimal local conditions for growth, with the global benefit of also regulating the global temperature. This means that the population size often reaches carrying capacity in the evolutionary Daisyworld when regulation occurs (e.g., Figure 10.3).

Regulation against small perturbations is performed by ecological interaction, since for a small change in external forcing it is possible for changes in population sizes of existing species to maintain the global temperature at the optimal level. This creates a homeostatic adaptive system where the regulatory mechanism is purely ecological. Evolutionary change does not play a part in regulation against very small perturbations, since at a fixed level of solar luminosity there will typically be a small number of dominant species which are well suited to current conditions and competitive exclusion prevents other species from gaining a foothold in the world.

However, if solar luminosity changes beyond a certain point, the current set of daisy species is unable to keep global temperature at the optimum for growth by purely ecolog-

ical means and homeostasis is lost. The current species are no longer so well suited to the conditions and become less competitive, creating an opportunity for new daisy species. New species are being constantly created by mutation, and if a variant occurs that is well matched to the current conditions it will rapidly multiply. Since such a species would necessarily have an effect that brought global temperature towards the optimum for growth, this will restore homeostasis.

The ecological regulatory mechanism is a form of elastic homeostatic adaptation for the same reasons that the two-species Daisyworld model was; changes in population sizes do not persist. The evolutionary regulatory mechanism, however, is a plastic homeostatic adaptive mechanism, since changes in the species composition of the Daisyworld are persistent until over-written by further evolutionary change. Changes in species composition also display hysteresis, since the current species composition depends not only on the current conditions but also on the recent history of the system. Also, since the evolutionary mechanism changes the species which constitute the ecological mechanism, the evolutionary Daisyworld displays second order homeostasis. The evolutionary homeostatic mechanism acts on the parameters of the ecological mechanism when the ecological mechanism alone is unable to regulate the global temperature. The evolutionary Daisyworld thus fits the scheme shown in Figure 14.5.

This homeostatic adaptive system becomes a homeorhetic adaptive system in the situation where regulation occurs but where there is variety in the preferred temperature for growth, i.e., when there are two possible growth functions (Figures 10.7, 10.8 and 10.9) or when preferred growth temperature adapts more slowly than albedo (Figure 10.10). Here regulation at any particular point in time occurs around a single growth function (with a single preferred temperature), since it is easier to use albedo to regulate temperature than to adapt the growth function to current conditions. The target of regulation may change however, if a change in external forcing occurs that is large enough to make the current target growth function un-competitive. At this point there will be a rapid switch to a new growth function and regulation around a new preferred growth temperature.

The homeorhetic adaptive system will display short periods of homeostatic adaptation around each preferred temperature, which over longer timescales gives a stepped progression in global temperature (as seen in Figures 10.7, 10.8, 10.9 and 10.10). This process is classified as homeorhetic adaptation because the set point around which regulation occurs is subject to change.

14.5.2 Flask model

The Flask model consists of an array of flasks, each containing a microbial community and receiving a steady chemical flux. The effect of the microbial community on the abiotic environment of a flask determines the growth rate of the microbes, creating a bi-directional feedback between the biota and their environment. Multi-level selection enabled by the spatial structure of the model may lead to the evolution of flask communities that regulate their environment to the optimal conditions for microbial growth.

It has yet to be demonstrated that the Flask model displays homeostasis, and for this reason we are unable to analyse it as a homeostatic adaptive system, but the Flask model as described contains the necessary ingredients for homeostatic adaptation. The essential variables of the Flask system are the vectors specifying the composition of the abiotic environments in each flask, and there are target ranges for these defined by the ranges within which microbes can grow and reproduce. There are a variety of adaptive mechanisms in the model, which include ecological dynamics within and between flasks, and selection pressure at the individual level and at the level of the flask community. Perturbations to the essential variables come from within the Flask world in the form of genetic drift at the individual level within flasks, and from the invasion and colonisation dynamics between flasks.

However, any further speculation on the Flask world as a homeostatic adaptive system would be pure conjecture and we will stop at this affirmation of the *necessary* conditions for homeostatic adaptation within the Flask model. Further analysis can be performed only when the Flask conditions have also been shown to be *sufficient* for homeostatic adaptation to occur.

14.5.3 The Earth system

Both Daisyworld and the Flask model are intended at some level to be models of how environmental homeostasis might be maintained by the biota in a manner consistent with evolutionary theory. However, by necessity and design both models are hugely simplified in comparison to the real Earth system, and this means that much caution must be exercised in seeking to draw conclusions from these models concerning the dynamics of the biosphere. Nonetheless, with this caveat in place, it can be argued that if Daisyworld and/or the Flask model are good models of the Earth system at some level, and are demonstrated to be homeostatic adaptive systems, then we should have the courage of our convictions and hypothesise that the Earth system also is a homeostatic adaptive system. It is certainly an adaptive system, by virtue of being the home for almost all known adaptive systems, but whether or not it can be shown to be homeostatic is an open

question.

14.6 Discussion

Having considered our case studies as homeostatic adaptive systems, we now try to build from these examples and make general observations concerning homeostatic adaptation. In this section we look at some aspects of homeostatic adaptive systems that allow them to be differentiated and defined; our discussion here applies similarly to homeorhetic adaptive systems.

14.6.1 Classification of perturbations

Perturbations to a homeostatic adaptive system can be grouped into two categories: those that cause a loss of homeostasis, and those that do not. A loss of homeostasis will trigger a response in the regulator, which will affect the behaviour of the reacting part. This change in behaviour may or may not result in the recovery of homeostasis, meaning that the class of perturbations that cause a loss of homeostasis can be sub-divided into those perturbations which cause a temporary loss of homeostasis that is later recovered by the action of the regulator, and those which cause a loss of homeostasis which is unrecoverable.

Let Q be the set of distinct parameter sets for the reacting part of a homeostatic adaptive system. Then for every parameterisation $q \in Q$ there is a set P_q of perturbations which do not cause a loss of homeostasis when the reacting part is configured with that parameterisation. With a set P of all possible perturbations, we can now define distinct subsets P_I , P_{II} and P_{III} that contain different sorts of perturbation. P_I is the set that contains all perturbations to the system that never cause a loss of homeostasis, P_{II} is the set of all perturbations that cause loss of homeostasis for some (but not all) parameterisations of the reacting part, and P_{III} is the set of perturbations that the system can therefore never adapt to since they cause loss of homeostasis for all parameterisations in Q . These sets are defined as follows:

$$\begin{aligned}
 P &= P_I \cup P_{II} \cup P_{III} \\
 P_I &= \bigcap_{q \in Q} P_q \\
 P_{II} &= \bigcup_{q \in Q} P_q - \bigcap_{q \in Q} P_q \\
 P_{III} &= \{P_I \cup P_{II}\}^c = \bigcap_{q \in Q} P_q^c
 \end{aligned}$$

For example, in the Homeostat some displacements of a magnet will be P_I perturbations that do not cause any units to lose homeostasis. Other perturbations will cause the magnets of one or more units to go outside the target range. If a uniselector setting is subsequently found such that circuit dynamics return all magnets to their target ranges, the perturbation is a P_{II} perturbation. If there is no setting of the uniselector which will bring all magnets back in bounds, the perturbation is a P_{III} perturbation. A further example of a P_{III} perturbation is one that leads to a fundamental disruption to the system, such as a blow from a sledgehammer. This kind of perturbation is outside the normal operating conditions of the Homeostat, and there is no feature of the Homeostat's machinery that can adapt to it.

14.6.2 Requisite variety

The variety of different parameterisations to the reacting part that the regulator can effect is key to homeostatic adaptation. The set of possible parameterisations of the reacting part determines the possible ways in which it can interact with the environment. The variety of parameterisations can be seen as the variety of modes of interaction with the world.

The number and types of perturbations a homeostatic adaptive system can adapt to is determined by the variety of parameterisations to the reacting part, and by the effect of different perturbations on the essential variables for each parameter set. Each parameterisation $q \in Q$ will generate an associated set of responses to perturbation in the reacting part. Some of these responses will lead to a homeostatic outcome, while others will lead to a non-homeostatic outcome. Every parameter set $q \in Q$ for the reacting part therefore has an associated set P_q of perturbations for which homeostasis is maintained, and a complementary set of perturbations P_q^c of perturbations that cause a loss of homeostasis. The process of homeostatic adaptation can be thought of as selecting a suitable parameter set q from Q so that homeostasis is maintained in the current conditions, i.e., so that the current perturbation is a member of P_q .

We can illustrate this concept with a simple example system. Suppose a homeostatic adaptive system exists that is subjected to a small discrete set of perturbations P , and has a reacting part which produces discrete responses that change dependent on its parameterisation from a set Q . Let $P = \{p1, p2, p3, p4, p5\}$ and $Q = \{q1, q2, q3\}$. Each parameter set in Q has a set of possible outcomes associated with the different perturbations in P . We can tabulate this scheme as shown in Table 14.1.

Looking at Table 14.1 we can see that the sets of perturbations which do not cause a loss of homeostasis for each parameter set are $P_{q1} = \{p1, p2, p3\}$, $P_{q2} = \{p1, p2, p4\}$ and $P_{q3} = \{p2, p3, p4\}$. We can also use the table to derive the set of perturbations which never cause a loss of homeostasis as $P_I = \{p2\}$, the set of perturbations which cause a loss

	q1	q2	q3
p1	O	O	X
p2	O	O	O
p3	O	X	O
p4	X	O	O
p5	X	X	X

Table 14.1: Outcomes resulting from perturbations p for different parameter sets q for the reacting part in an example homeostatic adaptive system. O denotes homeostatic outcome, X denotes non-homeostatic outcome. Perturbation $p2$ never causes a loss of homeostasis, while perturbation $p5$ always causes a loss of homeostasis. Perturbations $p1$, $p3$ and $p4$ all cause a loss of homeostasis with some parameter sets but not with others; the system can adapt to these perturbations if they cause a loss of homeostasis.

of homeostasis with some (but not all) parameter sets as $P_{II} = \{p1, p3, p4\}$, and the set of perturbations which always cause a loss of homeostasis as $P_{III} = \{p5\}$. We can observe that these sets accord with the rules for the membership of P_I , P_{II} and P_{III} that were stated above in Section 14.6.1.

Successful maintenance of homeostasis in a homeostatic system depends on sufficient variety in the responses of the regulator, in accordance with Ashby's Law of Requisite Variety [6]. For every perturbation, the interaction of the reacting part with its environment and the perturbation will determine the values of the essential variables. For homeostasis to be maintained against a perturbation there must be a mode of interaction for the reacting part which leads to a homeostatic outcome for that perturbation. The variety of modes of interaction is determined by the variety of responses of the regulatory effector mechanism, so that variety in the response of the effector determines the variety of perturbations that can be adapted to. The example given above uses discrete perturbations and responses, but the same principle applies equally to homeostatic mechanisms involving continuously varying perturbations and responses.

14.6.3 Relation between first and second order systems

Second order homeostasis acts on the parameters to a first order homeostatic regulator, which in turn acts on the parameters of the reacting part of a homeostatic system. This means that the variety of perturbations which a second order system can adapt to depends not only on the variety of responses in the first order regulator, but also on the variety in the second order regulator.

We can illustrate this concept with another simple example system. Retaining our notation from the previous section, suppose a second order homeostatic adaptive system exists that is subject to a set of perturbations P . Let the reacting part responses depend on the its parameter set q , which is selected from the set of possibilities Q by the first

(a) First order regulator parameterisation Q1				(b) First order regulator parameterisation Q2		
	q11	q12	q13		q21	q22
p1	O	O	X	p1	X	X
p2	X	X	X	p2	X	O
p3	O	X	O	p3	O	X
p4	X	X	X	p4	X	X

Table 14.2: Outcomes resulting from perturbations p for different parameter sets q for the reacting part in an example second order homeostatic adaptive system. O denotes homeostatic outcome, X denotes non-homeostatic outcome. There is one table of outcomes for each of the first order regulator parameter sets $Q1$ and $Q2$. Perturbations $p1, p2$ and $p3$ can be adapted to by the combined effects of the first and second order regulators, since there is an entry in at least one of the tables that gives a homeostatic outcome for these perturbations. The system can never adapt to perturbation $p4$, since neither parameterisation $Q1$ or $Q2$ allows the first order regulator to adapt to this perturbation.

order regulator. Furthermore, let the set Q be selected by the second order regulator from another set R of parameterisations to the first order regulator. Let $P = \{p1, p2, p3, p4\}$, $R = \{Q1, Q2\}$, $Q1 = \{q11, q12, q13\}$ and $Q2 = \{q21, q22\}$. We can tabulate this scheme as shown in Table 14.2.

Looking at Table 14.2 we can see that the sets of perturbations which can be adapted to by the first order regulator in each of its parameterisations $Q1$ and $Q2$ are $P_{Q1} = \{p1, p3\}$ and $P_{Q2} = \{p2, p3\}$ respectively. We can see from the table that for the whole system, all of the perturbations cause a loss of homeostasis in at least one of $Q1$ or $Q2$, so $P_I = \emptyset$, the empty set. The set of perturbations that cause a temporary loss of homeostasis but can be adapted to is $P_{II} = \{p1, p2, p3\}$ since there is a response in at least one of $Q1$ or $Q2$ that gives a homeostatic outcome for these perturbations. Table 14.2 also gives $P_{III} = \{p4\}$ since there is no response in either $Q1$ or $Q2$ that gives a homeostatic outcome for perturbation $p4$. We can observe that these sets accord with the rules for the membership of P_I , P_{II} and P_{III} that were stated above in Section 14.6.1.

Homeostatic adaptation in the second order homeostatic adaptive system involves the selection of an appropriate reacting part response by the first order regulator, which may in turn require a change of first order parameter set by the second order regulator. In our example, the first order regulator cannot adapt to perturbation $p2$ when in parameterisation $Q1$ since none of $q11, q12, q13$ gives a homeostatic outcome to $p2$. In this situation, the second order regulator would re-parameterise the first order regulator with parameterisation $Q2$, allowing the first order regulator to choose $q22$, which gives a homeostatic outcome.

14.6.4 Parameter selection mechanism

It is not sufficient only for Q to contain a parameter set that leads to the recovery of homeostasis, since there must also be a mechanism for selecting that value. In the Homeostat this mechanism is random selection, whereas in the homeostatic plastic neural networks parameters are changed according to a continuous directed mechanism. In the Daisy-world models, population size changes according to a growth equation, while evolutionary change occurs as random mutations.

Different mechanisms for changing the parameters of the reacting part will have different implications for adaptation. Random selection of new values is likely to be slow but has the advantage that it will eventually find a suitable parameter set if one exists. Directed mechanisms are likely to be quicker, but have a disadvantage in that they may make certain parameter sets inaccessible at certain times, reducing the scope for adaptation.

14.6.5 Organisation death and essential variables

Ashby's definition of essential variables considers them as the variables related to survival; if homeostasis is lost in an essential variable then the system cannot be said to survive in its original state. In the Homeostat, the essential variables are defined as the angular deviations of the magnets, with a target range of $\pm 45^\circ$. However, there is no concept of survival in the Homeostat, since the Homeostat cannot die. This raises the question of what the essential variables in the Homeostat really signify.

If the system is defined as having certain essential variables and then a perturbation causes a loss of homeostasis in these essential variables, then the system so defined ceases to exist. In the case of the Homeostat, the system that 'dies' when this occurs is a particular form of organisation; it is the Homeostat with its current uniselector settings. A loss of homeostasis triggers a 'death of organisation', but the homeostatic adaptive mechanism immediately causes the 'birth' of a new form of organisation.

The physical form of the Homeostat can be seen as the substrate in which a succession of different organisations are created and then destroyed. Each form of organisation (each instantiation of the reacting part with a particular set of parameters) survives until such a time as a perturbation causes a loss of homeostasis in essential variables (pushes a magnet out of bounds), at which point a new form is created. These forms may survive only fleetingly, for instance if a new form is created which does not restore homeostasis. Figure 14.6 shows organisation death and succession as a schematic diagram.

The idea of organisation death resulting from a loss of homeostasis in a plastic homeostatic adaptive system is general, and does not apply just to the Homeostat. The Homeostat was chosen as an illustrative example because of the clear distinction between different

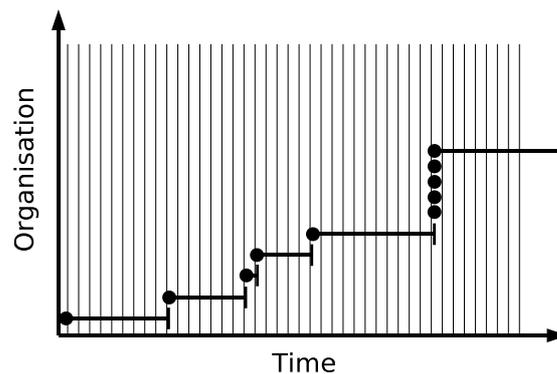


Figure 14.6: Organisation death and succession in the Homeostat. Different Homeostat organisations (different parameterisations) are arbitrarily compressed to a single dimension, as are different perturbations. Each circle represents the creation of a new organisation by the choosing of a new uniselector setting. This organisation persists for a time shown by the attached horizontal line. The vertical lines represent perturbations to the Homeostat, which sometimes cause a loss of homeostasis (short vertical tail to horizontal line). A loss of homeostasis represents ‘organisation death’ and is succeeded by a new organisation. Organisations can be short-lived or long-lived depending on the perturbations they experience.

forms that results from its discrete step-changes in parameters, but the concept also applies to continuous parameter changes.

14.6.6 Learning as second order homeostasis

Organisms are homeostatic adaptive systems. The concept of homeostasis was originally developed to describe the set of processes by which living things regulated their internal environment so that they continued in a living form. At the start of this thesis, in Chapter 1, we looked briefly at physiological homeostatic mechanisms for temperature and blood glucose regulation, which by our definition are elastic homeostatic adaptive systems.

The active behaviour of an organism can contribute to the maintenance of homeostasis. Animals, birds, fish, insects and various micro-organisms all interact purposefully with their environment in order to satisfy needs such as hunger or thirst, i.e., to regulate certain internal variables. We can consider a fixed pattern of behaviour as a first order elastic homeostatic adaptive system.

Some kinds of organism, such as mammals and birds, can also modify their pattern of interaction with the environment by learning. These organisms make long term and persistent changes to the way they behave, in addition to the short term changes that occur as they attend to different needs. Learning can be seen as the process by which the organisation of the animal changes from a non-adaptive to an adaptive state. This process involves successive changes in organisation, each involving the ‘death’ of one form of behaviour and the ‘birth’ of another. Each change is caused by a loss of homeostasis in

one or more of the essential variables associated with the behaviour-producing part, but this does not imply a loss of homeostasis in the essential variables associated with life. Through separation of the essential variables of the behaviour-producing subsystem from the essential variables of life, an animal may learn by a process of plastic homeostatic adaptation, without dying in the attempt.

Since the behaviour-producing subsystem of an animal is implicated in homeostatic regulation of essential variables related to the survival of the organism as a living system, and since learning acts on the parameters of the behaviour-producing subsystem, learning is a form of second order plastic homeostatic adaptation (cf. Ashby [7] and Di Paolo [46, 48]). The essential variables that are involved in the learning mechanisms can be the same as those which determine the overall survival of the organism, since the necessary separation between a loss of homeostasis triggering learning and a loss of homeostasis meaning death can be achieved by making the homeostatic range for the essential variables for the learning mechanism wider than that for survival.

14.6.7 Evolution as second order homeostasis

Here we argue that in a similar way to how learning can be viewed as a form of homeostatic adaptation [7, 46, 48], evolution can also be seen as a special case of homeostatic adaptation. The survival of an organism depends on the maintenance of homeostasis in the variables essential for life; failure to do so means death. The physiology and behaviour of the organism can be seen as a system for maintaining homeostasis in the face of perturbations that are either generated internally (such as hunger, thirst or old age) or come from the external environment (such as extreme conditions or predation). The organism (and its form of organisation) will survive as long as it can accommodate perturbations without losing homeostasis. This view of an organism as a first order homeostatic adaptive system implies a view of evolution as a second order adaptive mechanism acting on the parameters of the first order regulator.

Here the organism is seen as a homeostatic system with a fixed organisation determined by its genetic code. Lifetime learning is seen as a part of the homeostatic organisation, since it does not alter the genetic code. Reproduction of an organism also reproduces its genetic code and hence its organisation, i.e., the mode of an organism's interaction with the world is reproduced in its offspring, along with its physical form. Mutation and recombination play the role of mechanisms for altering the parameters of the organisation from generation to generation. Some forms of organisation will be better at withstanding perturbations than others and will therefore survive longer. Also, some forms of organisation will be better than others at reproducing themselves. We would expect forms of

organisation that are good at surviving and good at reproducing to proliferate.

Selection pressure takes the role of the second order homeostatic adaptive mechanism, allowing variations of the first order regulator that maintain homeostasis to prosper and removing variants that do not. Reproduction adds new individuals to the population and (by mutation) suggests new variants of the first order system. Implicitly, this considers a population as a vehicle for perpetuating a certain kind of homeostatic organisation and adapting that organisation to changes in the environment. These changes (which may be due to environmental or coevolutionary change) are experienced as new kinds of perturbation to the first order homeostatic system, which will cause some of the variants of the organisation to fail but allow others to maintain homeostasis and continue.

Tracing any lineage through evolutionary history gives a single-strand progression that looks very similar to the progression of the Homeostat shown in Figure 14.6. The process of evolutionary adaptation can be seen as a process of improving the organisation of the homeostatic system so that it can accommodate a greater range of perturbations from the environment. If we make the assumption that a lineage of successive organisational forms can be seen as repeated instantiations of different parameterisations of the same underlying structure, we have a process of temporally extended homeostatic adaptation. The homeostatic adaptor is not an identifiable physical entity (like the uniselector in the Homeostat) but it is instead an iterative process of variation and selection.

The view of evolution described above can be likened to the process by which the Homeostat deals with perturbations, but one in which many 'Homeostats' adapt in parallel. Where the Homeostat is a substrate in which different forms of organisation are serially created and destroyed, a population is a substrate that allows many forms of a certain kind of organisation to adapt in parallel. The evolution of the population is similar to a parallelised version of the trial-and-error adaptation of the Homeostat.

14.7 Conclusion

In this chapter we have considered the case study systems presented in this thesis as homeostatic adaptive systems. We have found that both the original Homeostat and the simulated Homeostat are first order plastic homeostatic adaptive systems, as are homeostatic plastic CTRNNs. The two-species Daisyworld model is a first order elastic homeostatic adaptive system. The evolutionary Daisyworld model is a second order homeostatic adaptive system, in which evolutionary adaptation performs plastic adaptation on the elastic first order ecological regulatory mechanism. More speculatively, we have suggested that the Flask model and the Earth system contain the necessary components to be homeostatic adaptive systems, but pointed out that these systems have yet to be shown to display

homeostasis.

In the latter section of the chapter we have looked at homeostatic adaptive systems more generally. It was noted homeostatic adaptation is a process of selecting an appropriate parameterisation for the reacting part so that the current perturbation does not disturb homeostasis, and that the variety in both the reacting part and the regulatory mechanism determine which perturbations can be adapted to and which cannot. The interpretation of essential variables in terms of system survival was discussed, and we argued that homeostatic adaptation involves ‘organisation death’ triggered by loss of homeostasis in the target variables. This view led to a way of thinking about biological learning and evolution as forms of second order homeostatic adaptation.

The stated intention of this chapter (and the thesis as a whole) was to work towards a ‘general theory of homeostatic adaptive systems’. Here we have presented a framework for categorising different kinds of homeostatic adaptive system, which we have applied to the case study systems covered in the earlier parts of the thesis. However, the question remains of whether we have made any progress towards constructing a unified theory. We have proposed a method of characterising homeostatic adaptive systems that allows us to construct a taxonomy of the different forms that homeostatic adaptation can take, and this is a valuable first step in the construction of a grander theory. However, good theories make predictions, and the framework for characterisation of homeostatic systems that we have proposed cannot be said to have any predictive power; for this reason we feel that we have fallen short of creating a general theory.

The utility of considering certain kinds of systems as homeostatic adaptive systems lies in the way it can guide our exploration of their behaviour. Where we observe constancy in the face of perturbation, we should look for regulation and homeostasis. Where we see homeostasis, we should look for homeostatic adaptation. If we can ascribe different parts of an observed system to different roles in the homeostatic adaptive system framework, we can move forward to a better understanding of system behaviour.

Chapter 15

Conclusion

15.1 Overview

In this final chapter we briefly summarise the material of Parts I to V, before drawing some final conclusions from the thesis as a whole.

15.2 Part I: Introduction

Chapter 1 began with an anecdote about different ways of regulating temperature in the home, which led into a statement of the main theme of this thesis: homeostatic adaptive networks. The concept of homeostasis was introduced with a brief history of the idea and a few examples of homeostasis from physiology, ecology and genetics. The idea of homeostasis as an organising principle that could play a major role in the determination of system behaviour was then discussed. A contrast was made between the traditional ‘adjective’ view of homeostasis and the ‘verb’ viewpoint argued for in this thesis. Here we treat homeostasis as a dynamic process shaping system development.

Different forms of homeostatic adaptation were identified: elastic homeostatic adaptation (where the system returns to its original state after adapting to perturbation), and plastic homeostatic adaptation (where the changes made during adaptation are persistent). After this informal discussion, some more rigorous definitions were given for homeostasis, elastic and plastic homeostatic adaptation. More complex schemes were described, such as systems with multiple homeostatic feedbacks, or systems with second order home-

ostatic adaptive mechanisms. These definitions were illustrated with examples from a toy model of a minimal homeostatic system.

The research aims of the thesis were then identified. These were the study of three homeostatic systems as case studies, followed by an abstraction of key properties from the case studies to build a general theory of plastic homeostatic adaptation. The three case study systems were briefly introduced: the Homeostat, homeostatic plasticity in neural networks, and environmental regulation by the biota. The intention was stated to study each system as an independent piece of research, before attempting to compare and contrast the different systems in the final section of the thesis.

Part I concluded with a few comments on simulation modelling methodology and a plan of the thesis layout.

15.3 Part II: The Homeostat

The first case study system was the Homeostat. The original Homeostat was an electromechanical device constructed by Ashby in the 1950s to demonstrate the principle of ultrastability. Ashby wanted to explain learning and adaptation in mammals, and suggested a theory where the target of learning was homeostasis of essential variables related to survival. Such a system could be created if the system was ultrastable to begin with; if its stability required all essential variables to be in homeostasis and exploratory parameter changes were triggered when homeostasis failed.

In Chapter 2, we developed a simulated Homeostat that was different to the original Homeostat in mechanism, but which operated along similar principles. The simulated Homeostat consisted of a number of units connected in a fully connected network architecture, with the activation of each unit determined by input from the other units and itself. If activation went outside a prescribed range, random change was triggered in afferent connection strengths and transfer function parameters. The simulated Homeostat was ultrastable, reliably converging to stability with all nodes in bounds. We showed that the time taken for convergence to stability after initialisation varied inversely with the number of units and with the tightness of the target range for activation. We also showed that the simulated Homeostat could adapt to perturbations in the form of external inputs to the units and that it could be used for simple control tasks.

15.4 Part III: Homeostatic plasticity in neural networks

Part III of the thesis described the development and application of homeostatic plastic mechanisms in neural networks for robot control. The aims were to improve evolutionary and neural robotics methodology by determining the properties of this novel mechanism, and if possible to gain insight into the role of neural homeostasis in biological nervous systems.

Chapter 3 gave the necessary background for this study. Evolutionary and neural robotics were introduced, with particular emphasis on continuous-time recurrent neural networks (CTRNNs) for robot control. Some issues with the evolutionary design of CTRNN controllers were discussed, in particular the problem of node saturation. Node saturation occurs when the range of input to a neuron is either too high or too low, leading to hyper-excitation or quiescence respectively. Node saturation means that much of CTRNN-space contains inert and unresponsive networks that are not suitable for use as controllers.

Chapter 3 also gives a brief overview of homeostatic plasticity in biological nervous systems. Homeostatic plasticity has recently been identified as an important mechanism in the brain, and involves the regulation of neural activity by a variety of mechanisms. The overall functional effects of homeostatic plasticity are uncertain, but it is thought to play a role in counter-balancing the positive feedbacks associated with Hebbian learning.

Prior to the work presented in this thesis there have been few attempts to incorporate homeostatic mechanisms into artificial neural networks, with those that exist either being based around Ashbyan homeostatic adaptation or being flawed in their execution. There is thus a useful contribution to be made by a principled study of the effects of biologically inspired homeostatic plastic mechanisms in artificial neural networks.

In Chapter 4 we developed homeostatic plastic mechanisms for use in CTRNNs. First of all, we identified the generic properties of the biological homeostatic plastic mechanisms most suited for implementation in CTRNNs. Then we derived simple mathematical forms for two different plasticity rules, one which implemented synaptic scaling and another based on an adaptive bias term. These rules were applied to neuron parameters when firing activity was too high or too low.

The utility of these mechanisms for preventing node saturation was explored in Chapter 5. The nature of the constraints imposed on the CTRNN by the inclusion of homeostatic plasticity was explored analytically, before some simple experiments were performed to demonstrate that homeostatic plasticity will always move a CTRNN away from saturation and towards a non-saturated constraint-satisfying region of parameter space (though external perturbations and initialisation conditions might prevent this region from

being reached).

The effects of homeostatic plasticity on network behaviour were examined in Chapter 6, which looked for effects at the levels of node, network and agent. It was found that homeostatic plasticity makes nodes and networks more sensitive to input, and allows signals to propagate much further through a network. At the agent level this makes the coupling between agent and environment much stronger, which was demonstrated using a simple simulated photo-sensitive agent. Homeostatic plasticity also makes networks more likely to display oscillatory dynamics.

It was noted in Chapter 6 that adding homeostatic plasticity to a CTRNN makes a qualitative change to the dynamics of the system. While homeostatic plasticity used purely as a developmental mechanism (i.e., applied for a period and then switched off) has the effects described in the preceding paragraph, leaving the homeostatic plastic mechanisms active adds new properties to the network. Slow oscillations are possible in a single node (when the plasticity is unable to find a stable attractor with activation in bounds and continually overshoots), and a new slower timescale is added to network dynamics.

Chapter 7 studied the utility of homeostatic plasticity for improving the evolvability of CTRNNs for robot control. A series of experiments was performed to test whether the application of homeostatic plasticity to CTRNNs would lead to more rapid or more reliable evolution of good robot controllers. A test platform was developed in which a simulated agent oriented its movement using an array of distance sensors in order to successfully catch a falling object or discriminate between different shapes. CTRNN controllers were evolved with homeostatic plasticity used either as a developmental mechanism, or as a continually active mechanism. Results were compared to the performance of standard non-plastic CTRNN controllers.

The results of the evolutionary experiments were inconclusive. When homeostatic plasticity was used purely as a developmental mechanism, a big improvement was observed in the speed of evolution of good controllers for one of the tasks, and a lesser improvement on the other task. When the plasticity was left active, the speed and consistency of evolution were improved, but there was a reduction in the ultimate level of performance achieved. A comparison with the evolutionary performance of centre-crossing CTRNN controllers suggested that the improvement offered by homeostatic plasticity stemmed largely from its ability to move networks towards the behaviour-rich region of parameter space around the centre-crossing condition.

The mixed set of results suggested that more evolutionary trials were needed to conclusively establish the utility of homeostatic plasticity for improving evolvability. The properties of increased sensitivity, improved signal propagation, and increased likelihood of oscillation, are often thought to be useful in adaptive behaviour, making the incon-

clusive nature of the evolvability results puzzling. It was concluded that the static task environment in the evolvability trials performed had favoured the non-plastic CTRNNs over the networks with plasticity, and that homeostatic plasticity might be most useful in situations where the environment or task parameters were subject to perturbation. Plastic networks offer a greater optimisation challenge than non-plastic networks due to the increased complexity of their dynamics. The static task environment used for the evolvability testing did not call for online adaptation and thus non-plastic controllers were easier to evolve. However, since homeostatic plasticity actively maintains network sensitivity and signal propagation during the network lifetime, it might be very useful for allowing a robot to operate in a variety of different environments or maintaining function in the face of perturbations, situations where traditional non-plastic networks would fail.

The implications of this work for neuroscience were also discussed in Chapter 8. It was acknowledged that the neural and plastic models used here are highly abstracted from their biological inspiration, so that extreme caution must be used when trying to draw conclusions from these simulations concerning biological nervous systems. With these caveats acknowledged, it was then argued that our findings support the view in neuroscience that homeostatic plasticity could play a role in counter-balancing the effects of Hebbian learning. However, it was also suggested that the results presented here contradict the view in neuroscience that homeostatic plasticity has a stabilising effect on network dynamics. Some neuroscientists suggest that the regulation of activity at individual nodes will also stabilise activity at the network level, but our results suggest the opposite; when the activity of individual nodes is regulated by homeostatic plasticity, network dynamics actually become less stable. Since homeostatic plasticity prevents node saturation, each node is made more sensitive to input and is forced to be active in network dynamics, making the network more likely to propagate signals and display oscillations, and raising overall levels of correlated activity.

15.5 Part IV: Homeostatic regulation of the environment by the biota

Part IV of the thesis is a study of homeostatic regulation of the environment by the biota. It is largely centred on Gaia theory, the idea that life regulates the biosphere to conditions suitable for life, but it is intended that more general implications should be drawn.

Chapter 9 presents a review of Gaia theory, from its inception as the controversial Gaia hypothesis, through to its current status as a valid topic for scientific research. Gaia theory rests on the existence of two-way feedback between the biota and its physical envi-

ronment, and states that the collective effect of individual-level selection for environment-altering traits is global regulation. The main area of controversy over Gaia theory concerns its compatibility with Darwinian evolution. Supporters claim that the continued existence of a habitable environment on Earth in the face of perturbations (such as changes in solar luminosity and asteroid impacts) is strong evidence for the presence of regulatory feedbacks between life and its environment. They also claim that global regulation can emerge from individual-level selection without any need for conscious control. Critics of the Gaia theory claim that it would be vulnerable to ‘cheaters’, that destabilising feedbacks are just as likely to occur as regulatory ones, and that the continued presence of life on Earth is nothing more than a criterion for the evolution of sentient observers.

Chapter 9 moves on to describe the Daisyworld model, giving details of the original model and describing the various extensions to it that have appeared in the literature since then. Daisyworld displays regulation of global temperature through ecological competition between black and white daisy species, in the face of increasing solar insolation. The Guild model is also described, which is based on similar assumptions to Daisyworld, but includes the formation of regulating ecologies.

Chapter 9 finishes with a statement of the research aims for Part IV, which are defined as gaining a full understanding of the conditions under which regulation will evolve in Daisyworld and developing the Flask model, a new model which relaxes the assumptions of Daisyworld and allows for the evolution of environmental regulation by multi-level selection in microbial communities.

A simplified Daisyworld model is developed in Chapter 10, in which the physics is stripped down and an individual-based approach is used to facilitate the easy inclusion of Darwinian evolution. The simplified Daisyworld model is first used to replicate known results from more conventional Daisyworld models, such as ecological regulation of temperature by competition between black and white daisies, and the failure of regulation when the preferred temperature for growth is allowed to mutate freely. Next the model is used to explore the importance of constraints on adaptation for the evolution of regulation. The main finding is that constraints on adaptation of growth temperature are essential for regulation to occur. If there are no constraints on adaptation of preferred growth temperature, daisies simply evolve to prefer the current ambient temperature and do not use albedo to moderate their local environment. When two possible preferred growth temperatures coexist, competitive exclusion leads to regulation around one of them. When adaptation of growth function is unconstrained in range, but occurs much more slowly than adaptation of albedo, regulatory epochs occur in which competitive exclusion causes regulation around a particular growth temperature until it is made non-viable by changes in external forcing.

We also found that regulation fails when two of the core assumptions of the Daisyworld model are relaxed. If daisies cannot create a local buffer against the environment, or if selfish adaptations are permitted to have a non-regulatory environmental impact, regulation of global temperature does not occur. These assumptions also apply to the Guild model, and it is argued that a new model is needed to study the possibility of environmental regulation in an evolving ecosystem where these assumptions do not hold.

The Flask model is introduced in Chapter 11 in the form of a sketch of its logical structure and the assumptions it is based on. The Flask model consists of microbial communities inhabiting an array of flasks, each of which is supplied with a flow of nutrients. Competition, reproduction, and mutation of microbes allows the possible formation of stable ecologies in each flask. Microbes have an effect on their abiotic environment, which in turn has an effect on growth, creating a double feedback between a flask community and its environment. Limited migration between flasks allows for colonisation of neighbouring flasks by successful microbial communities, and this allows the possibility of higher-level selection on the interaction between a community and its environment. The main hypothesis of the Flask model is that multi-level selection (enabled by the spatial structure of the Flask world) can select for communities that improve their environment; the global effect of this will be environmental regulation.

The assumptions that the Flask model is based on are discussed at the end of Chapter 11, and in Chapter 12 some experimental modelling work is described that seeks to establish the validity of one of these assumptions. An experiment is described in which flask communities are artificially selected for the composition of their abiotic environment. Over time, a significant response to selection is observed, demonstrating that heritability can exist between a parent flask community and its colonist offspring. The size of the response to selection is inversely proportional to the amount of time for which each flask community is allowed to develop before selection occurs.

The implications of the results presented in Part IV are discussed in Chapter 13. It is argued that the simplified Daisyworld model is useful because it allows clear understanding of the necessary conditions for regulation to occur, and also because its simplicity makes it more general than the original Daisyworld model. The new results concerning competitive exclusion and regulatory epochs are also novel to the Daisyworld literature, as is the recognition of the fundamental importance of the Daisyworld assumptions of a local buffer and selfish adaptation contributing to global regulation.

It is suggested that the Flask model can have a significant contribution to the Gaia debate if it is developed to fruition, but that in its current status as a thought experiment its main contribution to Gaia theory is in the hypothesis it embodies. It is also possible that the Flask model can be useful for studying multi-level selection in evolutionary theory,

particularly as its basis in microbial ecology means that it creates hypotheses testable by laboratory experiments with real microbial communities.

15.6 Part V: Discussion

In Chapter 14 we drew together the different strands of research presented in the thesis and analysed each case study system in terms of homeostatic adaptation. We found that both the original and simulated Homeostats studied in Part II were first order plastic homeostatic adaptive systems, as were the homeostatic plastic CTRNNs studied in Part III. The models of environmental regulation by the biota we studied in Part IV had different properties. The two-species Daisyworld model was a first order elastic homeostatic adaptive system, while the evolutionary Daisyworld was a second order homeostatic adaptive system with both plastic and elastic adaptive mechanisms. The evolutionary Daisyworld could also be a homeorhetic adaptive system in the cases where multiple targets for regulation existed. The Flask model and Earth system were observed to contain the necessary components to be homeostatic adaptive systems, but it was noted that these systems have yet to be shown to be homeostatic.

The latter section of Chapter 14 considered some properties that are general to all homeostatic adaptive systems. Homeostatic adaptation was described as the process of selecting an appropriate response to perturbation so that homeostasis is maintained. Variety of response in both the reacting part and the regulator was identified as important for regulation, since this variety determines which perturbations can be adapted to. Some discussion of different classes of perturbation was given, noting that some perturbations never cause loss of homeostasis, while some perturbations always cause loss of homeostasis. The relationship between first and second order homeostasis was examined, before some discussion of the concept of ‘organisation death’, where a loss of homeostasis in a plastic homeostatic adaptive system means that the system will cease to exist in its current form of organisation. This led to some speculation that learning and evolution could be viewed as second order homeostatic adaptive systems.

The remainder of Part V consists of the summary of the thesis that is given in this chapter.

15.7 Conclusion

The aim of the thesis was to study homeostatic adaptive networks through the study of three case studies. This aim has been met; we have looked at the Homeostat, homeostatic

plasticity in neural networks, and environmental regulation by the biota, and drawn some general conclusions concerning homeostatic adaptation.

Of the case studies, the most important contributions are made in the studies of homeostatic plasticity in neural networks and environmental regulation by the biota. The work on the Homeostat is valid, and useful as a didactic tool to explain some of the concepts of homeostatic adaptation, but breaks little new scientific ground. In contrast, the implementation in artificial neural networks of homeostatic plastic mechanisms based on those found in biology is a novel contribution to the literature. The results from this section should be of interest to the evolutionary robotics community, and suggest new avenues for future research. Likewise, the work presented in simplifying and extending the Daisyworld is a useful contribution to the Gaia theory literature, and the development of the Flask model is an original contribution that may hope to answer the call for new models in the Gaian debate.

In hindsight, the wide remit of the thesis (i.e., the study of homeostatic adaptive networks) has been both liberating and problematic. While the topic of homeostasis allows for study of a variety of interesting systems, it has been difficult to study each of these systems in sufficient depth to do useful research without sacrificing the overall aim. Other homeostatic systems might have been studied in place of those chosen here, or a greater number of systems might have been studied in less detail. We hope to have struck a balance between the depth of research into each case study and the breadth of different case studies chosen, so that useful results were achieved from each individual piece of work as well as from the thesis as a whole. A major regret is that the Flask model could not be developed fully here, but as mentioned previously, time and space constraints did not allow this to be done.

Concerning the central theme of homeostatic adaptive networks, we do not pretend to have discovered any new phenomena or created a new theory. Instead we hope to have highlighted the significance of homeostatic adaptation as an organising principle in biology (and in other systems) and to have gone some way towards clarifying the different forms homeostatic adaptation can take. The utility of this work, as with most cybernetic research, lies less in the new theories and techniques it produces than in the way of thinking it puts forward. The view of systems as homeostatic adaptive networks is useful in the way it can aid understanding and suggest directions for future research.

Bibliography

- [1] L.F. Abbott and S.B. Nelson. Synaptic plasticity: Taming the beast. *Nature Neuroscience*, 3:1178–1183, 2000.
- [2] L.F. Abbott, G.G. Turrigiano, G. LeMasson, and E. Marder. Activity-dependent conductances in model and biological neurons. In D.L. Waltz, editor, *Natural and Artificial Parallel Computing*, pages 43–68, NY, 1996. SIAM.
- [3] G.J. Ackland, M.A. Clark, and T.M. Lenton. Catastrophic desert formation in Daisyworld. *Journal of Theoretical Biology*, 223(1):39–44, 2003.
- [4] American Geophysical Union. *Abstracts from the Second Chapman Conference on the Gaia Hypothesis, Valencia, Spain, 19-23 June, 2000*, 2000.
- [5] M.O. Andreae and P.J. Crutzen. Atmospheric aerosols: Biogeochemical sources and role in atmospheric chemistry. *Science*, 276:1052–1058, 1997.
- [6] W. R. Ashby. *Introduction to Cybernetics*. Chapman & Hall, London, 1956.
- [7] W.R. Ashby. *Design for a Brain*. Chapman and Hall, London, 1952.
- [8] A. Balaam. Analysing homeostatic adaptation to sensorimotor disruptions. Master’s thesis, School of Computing and Cognitive Sciences, University of Sussex, 2001.
- [9] L. Barnett. Tangled webs: Evolutionary dynamics on fitness landscapes with neutrality. Master’s thesis, School of Computing and Cognitive Sciences, University of Sussex, 1997.
- [10] R. Beer. Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes, M.J. Mataric, J-A. Meyer, J.B. Pollack, and S.W. Wilson, editors, *From Animals to Animats 4: Proceedings of the 4th Intl. Conf. on Simulation of Adaptive Behavior*, pages 421–429, Cambridge, MA, 1996. MIT Press.

- [11] R. Beer. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243, 2003.
- [12] R. Beer and J. Gallagher. Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1:91–122, 1992.
- [13] R.D. Beer. *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology*. Academic Press, 1990.
- [14] R.D. Beer. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215, 1995.
- [15] R.D. Beer. On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3:469–509, 1995.
- [16] R.D. Beer. Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3):91–99, 2000.
- [17] R.D. Beer, H.J. Chiel, and J.C. Gallagher. Evolution and analysis of model CPGs for walking II: General principles and individual variability. *Journal of Computational Neuroscience*, 7(2):119–147, 1999.
- [18] R.D. Beer, H.J. Chiel, R.D. Quinn, K. Espenschied, and P. Larsson. A distributed neural network architecture for hexapod robot locomotion. *Neural Computation*, 4(3):356–365, 1992.
- [19] M. Begon, J.L. Harper, and C.R. Townsend. *Ecology: Individuals, Populations and Communities (Third Edition)*. Blackwell Science, Oxford, 1996.
- [20] C. Bernard. *Leçons sur les propriétés physiologiques et les alterations pathologiques des liquides de l'organisme*. Bailliere, Paris, France, 1859.
- [21] S. Bovet, M. Fend, and R. Pfeifer. Simulating whisker sensors - on the role of material properties for morphology, behavior and evolution. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer, editors, *From Animals to Animats 8: Proceedings of the 8th Intl. Conf. on Simulation of Adaptive Behavior*, Cambridge, MA, 2004. MIT Press.
- [22] R.A. Brooks. Intelligence without reason. In J. Myopoulos and R. Reiter, editors, *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, Sydney, Australia, 1991. Morgan Kaufmann.

- [23] R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
- [24] C.L. Buckley, S. Bullock, and N. Cohen. Timescale and stability in adaptive behaviour. In M.S. Capcarrere, A.A. Freitas, P.J. Bentley, C.G. Johnson, and J. Timmis, editors, *Proceedings of the 8th European Conference on Artificial Life (ECAL2005)*, Berlin Heidelberg, 2005. Springer Verlag.
- [25] J. Burrone and V.N. Murthy. Synaptic gain control and homeostasis. *Current Opinion in Neurobiology*, 13:560–567, 2003.
- [26] W.B. Cannon. *The Wisdom of the Body*. Norton, New York, 1932.
- [27] R.J. Charlson, J.E. Lovelock, M.O. Andreae, and S.G. Warren. Oceanic phytoplankton, atmospheric sulphur, cloud albedo and climate. *Nature*, 326:655–661, 1987.
- [28] H.J. Chiel and R.D. Beer. The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20:553–557, 1997.
- [29] H.J. Chiel, R.D. Beer, and J.C. Gallagher. Evolution and analysis of model CPGs for walking I: Dynamical modules. *Journal of Computational Neuroscience*, 7(2):99–118, 1999.
- [30] A. Clark. *Being There: Putting Brain, Body and World Together Again*. MIT Press, Cambridge, MA, 1997.
- [31] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):71–104, 1993.
- [32] D. Cliff, P. Husbands, and I. Harvey. Evolving visually guided robots. In J-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2: Proc. of the Second Intl. Conf. on Simulation of Adaptive Behavior (SAB92)*, pages 374–383, Cambridge MA, 1992. MIT Press/Bradford Books.
- [33] D. Cliff and G.F. Miller. Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M.J. Mataric, J-A Meyer, J.B. Pollack, and S.W. Wilson, editors, *From Animals to Animats 4*, pages 506–515, Cambridge, MA, 1996. MIT Press.

- [34] D. T. Cliff. Computational neuroethology: A provisional manifesto. Technical Report CSRP 162, School of Cognitive and Computing Sciences (COGS), University of Sussex, May 1990.
- [35] J.E. Cohen and A.D. Rich. Interspecific competition affects temperature stability in Daisyworld. *Tellus*, 52B(3):980–984, 2000.
- [36] C. Darwin. *On the Origin of Species*. John Murray, London, 1859.
- [37] G.W. Davis and I. Bezprozvanny. Maintaining the stability of neural function: A homeostatic hypothesis. *Annual Review of Physiology*, 63:847–869, 2001.
- [38] G.W. Davis and C.S. Goodman. Genetic analysis of synaptic development and plasticity: Homeostatic regulation of synaptic efficacy. *Current Opinion in Neurobiology*, 8:149–156, 1998.
- [39] R. Dawkins. *The Selfish Gene*. Oxford University Press, New York, 1976.
- [40] R. Dawkins. *The Extended Phenotype*. Oxford University Press, Oxford, 1982.
- [41] S. De Gregorio, R. A. Pielke, and G.A. Dalu. A delayed biophysical system for the Earth’s climate. *Journal of Nonlinear Science*, 2:293–318, 1992.
- [42] S. De Gregorio, R. A. Pielke, and G.A. Dalu. Feedback between a simple biosystem and the temperature of the Earth. *Journal of Nonlinear Science*, 2:263–292, 1992.
- [43] D.C. Dennett. Cognitive wheels: The frame problem of AI. In C. Hookway, editor, *Minds, Machines, and Evolution: Philosophical Studies*, pages 129–151. Cambridge University Press, 1984.
- [44] N.S. Desai, L.C. Rutherford, and G.G. Turrigiano. BDNF regulates the intrinsic excitability of cortical neurons. *Learning and Memory*, 6:284–291, 1998.
- [45] N.S. Desai, L.C. Rutherford, and G.G. Turrigiano. Plasticity in the intrinsic excitability of neocortical pyramidal neurons. *Nature Neuroscience*, 2:515–520, 1999.
- [46] E. Di Paolo. Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In J-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 6: Proceedings of the 6th Intl. Conf. on Simulation of Adaptive Behavior*, pages 440–449, Cambridge, MA, 2000. MIT Press.

- [47] E. Di Paolo. Evolving robust robots using homeostatic oscillators. Technical Report CSRP 548, School of Cognitive and Computing Sciences, University of Sussex, UK, 2002.
- [48] E. Di Paolo. Organismically-inspired robotics: Homeostatic adaptation and natural teleology beyond the closed sensorimotor loop. In K. Murase and T. Asakura, editors, *Dynamical Systems Approach to Embodiment and Sociality*, pages 19–42, Adelaide, 2003. Advanced Knowledge International.
- [49] E.A. Di Paolo. Rhythmic and non-rhythmic attractors in asynchronous random Boolean networks. *BioSystems*, 59(3):185–195, 2001.
- [50] E.A. Di Paolo. Spike timing dependent plasticity for evolved robots. *Adaptive Behavior*, 10(3/4):243–263, 2002.
- [51] E.A. Di Paolo, J. Noble, and S. Bullock. Simulation models as opaque thought experiments. In M.A. Bedau, J.S. McCaskill, N.H. Packard, and S. Rasmussen, editors, *Artificial Life VII: The Seventh International Conference on the Simulation and Synthesis of Living Systems*, pages 497–506, Cambridge, MA, 2000. MIT Press/Bradford Books.
- [52] W.F. Doolittle. Is nature really motherly? *The Coevolution Quarterly*, Spring:58–63, 1981.
- [53] K. Downing. Euzone: Simulating the evolution of aquatic ecosystems. *Artificial Life*, 3(4):307–333, 1997.
- [54] K. Downing. The simulated emergence of distributed environmental control in evolving microcosms. *Artificial Life*, 6(2):123–153, 2002.
- [55] K. Downing and P. Zvirinsky. The simulated evolution of biochemical guilds: Reconciling Gaia theory and natural selection. *Artificial Life*, 5(4):291–318, 1999.
- [56] G. Dudek and M. Jenkin. *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, 2000.
- [57] P. Erdős and A. Renyi. On random graphs I. *Publicationes Mathematicae*, 6:290–297, 1959.
- [58] M. Fend, H. Yokoi, and R. Pfeifer. Optimal morphology of a biologically-inspired whisker array on an obstacle-avoiding robot. In W. Banzhaf, T. Christaller, P. Dittich, J.T. Kim, and J. Ziegler, editors, *Proceedings of the 7th European Conference on Artificial Life (ECAL)*. Springer, 2003.

- [59] D. Floreano and C. Mattiussi. Evolution of spiking neural controllers for autonomous vision-based robots. In T. Gomi, editor, *Evolutionary Robotics IV*. Springer-Verlag, Berlin, 2001.
- [60] D. Floreano and F. Mondada. Evolution of plastic neurocontrollers for situated agents. In P. Maes, M.J. Mataric, J-A. Meyer, J.B. Pollack, and S.W. Wilson, editors, *From Animals to Animats 4: Proceedings of the 4th Intl. Conf. on Simulation of Adaptive Behavior*, Cambridge, MA, 1996. MIT Press.
- [61] D. Floreano and J. Urzelai. Evolution of neural controllers with adaptive synapses and compact genetic encoding. In D. Floreano, J-D. Nicoud, and F. Mondada, editors, *Proceedings of the 5th European Conference on Artificial Life (ECAL)*, Berlin, 1999. Springer-Verlag.
- [62] D. Floreano and J. Urzelai. Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13:431–443, 2000.
- [63] D. Floreano and J. Urzelai. Evolution of plastic control networks. *Autonomous Robots*, 11:311–317, 2001.
- [64] D. Floreano and J. Urzelai. Neural morphogenesis, synaptic plasticity, and evolution. *Theory in Biosciences*, 120:225–240, 2001.
- [65] K.I. Funahashi and Y. Nakamura. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, 6:801–806, 1993.
- [66] J.C. Gallagher, R.D. Beer, K.S. Espenschied, and R.D. Quinn. Application of evolved locomotion controllers to a hexapod robot. *Robotics and Autonomous Systems*, 19:95–103, 1996.
- [67] L. Glass and M.C. Mackey. *From Clocks to Chaos: The Rhythms of Life*. Princeton University Press, Princeton, NJ, 1988.
- [68] D. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
- [69] J. Golowasch, L.F. Abbott, and E. Marder. Activity-dependent regulation of ionic conductances in an identified neuron of the stomatogastric ganglion of the crab *Cancer borealis*. *J. Neurosci.*, 19:RC33(1–5), 1999.
- [70] J. Golowasch, M. Casey, L.F. Abbott, and E. Marder. Network stability from activity-dependent regulation of neuronal conductances. *Neural Computation*, 11:1079–1096, 1999.

- [71] C.J. Goodnight and L. Stevens. Experimental studies of group selection: What do they tell us about group selection in nature. *American Naturalist*, 150:S59–S79, 1997.
- [72] F. Gruau. *Neural Network Synthesis using Cellular Encoding and the Genetic Algorithm*. PhD thesis, Laboratoire de l’Informatique du Parallélisme, Ecole Normale Supérieure de Lyon, France, 1994.
- [73] H. Haken. *Synergetics*. Springer, Berlin, 1983.
- [74] P. Hammerstein. Darwinian adaptation, population genetics and the streetcar theory of evolution. *Journal of Mathematical Biology*, 34(5-6):511–532, 1996.
- [75] S.P. Harding. Food web complexity enhances community stability and climate regulation in a geophysiological model. *Tellus*, 51B:815–829, 1999.
- [76] S.P. Harding and J. E. Lovelock. Exploiter-mediated coexistence and frequency-dependent selection in a numerical model of biodiversity. *Journal of Theoretical Biology*, 182:109–116, 1996.
- [77] S. Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.
- [78] I. Harvey. Evolutionary robotics and SAGA: the case for hill crawling and tournament selection. Technical Report CSRP222, School of Cognitive and Computing Sciences, University of Sussex, 1992.
- [79] I. Harvey. Artificial evolution and real robots. *Artificial Life and Robotics*, 1:35–38, 1997.
- [80] I. Harvey. Homeostasis and rein control: From Daisyworld to active perception. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R.A. Watson, editors, *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE’9*, pages 309–314, Cambridge, MA, 2004. MIT Press.
- [81] I. Harvey and T. Bossomaier. Time out of joint: Attractors in asynchronous random boolean networks. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 67–75, Cambridge, MA, 1997. MIT Press.
- [82] I. Harvey, P. Husbands, and D. Cliff. Issues in evolutionary robotics. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2: Proc. of the Second Intl. Conf. on Simulation of Adaptive Behavior (SAB92)*, pages 364–373, Cambridge MA, 1992. MIT Press/Bradford Books.

- [83] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, 1998.
- [84] T. Hoinville and P. Henaff. Evolving plastic neural controllers stabilized by homeostatic mechanisms for adaptation to a perturbation. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. A. Watson, editors, *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 81–87, Cambridge, MA, 2004. MIT Press.
- [85] J.H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI, 1975.
- [86] J. Hopfield. Neural networks and physical systems with emergent collective computational properties. *Proceedings of the National Academy of Sciences of the USA*, 79:2554–2588, 1982.
- [87] P. Husbands, T.M.C Smith, N. Jakobi, and M. O’Shea. Better living through chemistry: Evolving GasNets for robot control. *Connection Science (Special Issue: BioRobotics)*, 10(3-4):185–210, 1998.
- [88] A.J. Ijspeert, J. Hallam, and D. Willshaw. From lampreys to salamanders: evolving neural controllers for swimming and walking. In R. Pfeifer, B. Blumberg, J-A. Meyer, and S.W. Wilson, editors, *From Animals to Animats 5: Proceedings of the Fifth International Conference of The Society for Adaptive Behavior (SAB98)*, pages 390–399, Cambridge, MA, 1998. MIT Press.
- [89] N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6(2):325–368, 1997.
- [90] N. Jakobi. Half-baked, ad-hoc, and noisy: Minimal simulations for evolutionary robotics. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, 1997.
- [91] N. Jakobi. *Minimal Simulations for Evolutionary Robotics*. PhD thesis, School of Cognitive and Computing Sciences, University of Sussex, 1998.
- [92] N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. *Lecture Notes in Computer Science*, 929:704–??, 1995.
- [93] S.D. Jascourt and W. H. Raymond. Comments on “Chaos in Daisyworld” by X. Zeng et al. *Tellus*, 44B:243–246, 1992.

- [94] H. Jonas. *The phenomenon of life: Towards a philosophical biology*. Northwestern University Press, Chicago, 1966.
- [95] S.A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, 1993.
- [96] R. Keeling. Mechanisms for stabilization and destabilization of a simple biosphere: Catastrophe on Daisyworld. In S. H. Schneider and P. J. Boston, editors, *Scientists on Gaia: Papers delivered at the American Geophysical Union's annual Chapman Conference in March, 1988*, pages 118–120, London, 1991. MIT Press.
- [97] J.W. Kirchner. The Gaia hypotheses: Are they testable? Are they useful? In S.H. Schneider and P. J. Boston, editors, *Scientists on Gaia: Papers delivered at the American Geophysical Union's annual Chapman Conference in March, 1988*, Cambridge, MA:, 1991. MIT Press.
- [98] J.W. Kirchner. The Gaia hypothesis: Fact, theory, and wishful thinking. *Climatic Change*, 52:391–408, 2002.
- [99] J.W. Kirchner. The Gaia hypothesis: Conjectures and refutations. *Climatic Change*, 58:21–45, 2003.
- [100] H. Kitano. Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4(4):461–476, 1990.
- [101] A. Kleidon. Testing the effect of life on Earth's functioning: How Gaian is the Earth system? *Climatic Change*, 52:383–389, 2002.
- [102] K.N. Laland, F.J. Odling-Smee, and M.W. Feldman. On the evolutionary consequences of niche construction. *Journal of Evolutionary Biology*, 9:293–316, 1996.
- [103] K.N. Laland, F.J. Odling-Smee, and M.W. Feldman. Niche construction, biological evolution and cultural change. *Behavioral and Brain Sciences*, 23:131–175, 1999.
- [104] J.A. Larsen. *The Boreal Ecosystem*. Academic Press, New York, 1980.
- [105] T. Lenton. Gaia and natural selection. *Nature*, 394:439–447, 1998.
- [106] T.M. Lenton. Testing Gaia: The effect of life on Earth's habitability and regulation. *Climatic Change*, 52:409–422, 2002.
- [107] T.M. Lenton and J. E. Lovelock. Daisyworld is Darwinian: constraints on adaptation are important for planetary self-regulation. *Journal of Theoretical Biology*, 206:109–114, 2000.

- [108] T.M. Lenton and J.E. Lovelock. Daisyworld revisited: quantifying biological effects on planetary self-regulation. *Tellus*, 53B:288–305, 2001.
- [109] T.M. Lenton and M. van Oijen. Gaia as a complex adaptive system. *Phil. Trans. R. Soc. Lond. Series B*, 357:683–695, 2002.
- [110] T.M. Lenton and A. J. Watson. Redfield revisited 1: Regulation of nitrate, phosphate and oxygen in the ocean. *Global Biogeochemical Cycles*, 14(1):225–248, 2000.
- [111] T.M. Lenton and D.M. Wilkinson. Developing the Gaia theory: A response to the criticisms of Kirchner and Volk. *Climatic Change*, 58:1–12, 2003.
- [112] I.M. Lerner. *Genetic Homeostasis*. John Wiley, New York, 1954.
- [113] R.C. Lewontin. The units of selection. *Annual Review of Ecology and Systematics*, 1:1–18, 1970.
- [114] P.S. Liss, A.D. Hatton, G. Malin, P.D. Nightingale, and S.M. Turner. Marine sulphur emissions. *Phil. Trans. R. Soc. Lond. B*, 352:159–169, 1997.
- [115] A.J. Lotka. *Elements of physical biology*. Williams & Wilkins Co., Baltimore, 1925.
- [116] J.E. Lovelock. *Gaia*. Oxford University Press, Oxford, 1979.
- [117] J.E. Lovelock. Gaia as seen through the atmosphere. In P. Westbroek and E. W. de Jong, editors, *Biomineralization and Biological Metal Accumulation*, pages 15–25, Dordrecht, 1982. Reidel.
- [118] J.E. Lovelock. Daisy World - a cybernetic proof of the Gaia hypothesis. *The Coevolution Quarterly*, Summer:66–72, 1983.
- [119] J.E. Lovelock. Geophysiology. *Transactions of the Royal Society of Edinburgh: Earth Sciences*, 80:169–175, 1989.
- [120] J.E. Lovelock. Geophysiology - the science of Gaia. In S.H. Schneider and P. J. Boston, editors, *Scientists on Gaia: Papers delivered at the American Geophysical Union's annual Chapman Conference in March, 1988*, Cambridge, MA:, 1991. MIT Press.
- [121] J.E. Lovelock. A numerical model for biodiversity. *Philosophical Transactions of the Royal Society of London: Series B*, 338:383–391, 1992.

- [122] J.E. Lovelock. A geophysiologicalist's thoughts on the natural sulphur cycle. *Phil. Trans. R. Soc. B*, 352:143–147, 1997.
- [123] J.E. Lovelock and L.R. Kump. Failure of climate regulation in a geophysiological model. *Nature*, 369:732–734, 1994.
- [124] J.E. Lovelock and L. Margulis. Atmospheric homeostasis by and for the biosphere: the Gaia hypothesis. *Tellus*, 26(2):2–9, 1973.
- [125] L. Maddock. Effects of simple environmental feedback on some population models. *Tellus*, 43B:331–337, 1991.
- [126] L. Margulis. *The Symbiotic Planet: A New Look at Evolution*. Basic Books, New York, 1998.
- [127] B. Mathayomchan and R. D. Beer. Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation*, 14:2043–2051, 2002.
- [128] H. Maturana and F.J. Varela. *Autopoiesis and cognition: the realization of the living*. D. Reidel Publishing, Dordrecht, Holland, 1980.
- [129] H. Maturana and F.J. Varela. *The tree of knowledge: The biological roots of human understanding*. Shambhala, Boston, MA, 1987.
- [130] J. Maynard Smith. Group selection and kin selection. *Nature*, 201:1145–1147, 1964.
- [131] W. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. volume 5, pages 115–133. 1943.
- [132] O. Michel. Evolutionary neurogenesis applied to mobile robotics. In M. Patel and V. Honavar, editors, *Advances In Evolutionary Synthesis of Neural Systems*. MIT Press, Cambridge, MA, 1999.
- [133] M. Minsky and S.A. Papert. *Perceptrons*. MIT Press, Cambridge, MA, 1969.
- [134] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, 1996.
- [135] C. Nevison, V. Gupta, and L. Klinger. Self-sustained oscillations on Daisyworld. *Tellus*, 51B(4):806–814, 1999.
- [136] M.E.J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

- [137] A. Penn. Modelling artificial ecosystem selection: A preliminary investigation. In W. Banzhaf, T. Christaller, P. Dittrich, J.T. Kim, and J. Ziegler, editors, *Advances in Artificial Life, Proceedings of the 7th European Conference on Artificial Life (ECAL 2003)*, volume 2801 of *Lecture Notes in Computer Science*, pages 659–666. Springer, 2003.
- [138] A. Penn and I. Harvey. The role of non-genetic change in the heritability, variation, and response to selection of artificially selected ecosystems. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R.A. Watson, editors, *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE 9)*, pages 352–357, Cambridge MA, 2004. MIT Press.
- [139] J.B. Pollack, H. Lipson, G.S. Hornby, and P. Funes. Three generations of automatically designed robots. *Artificial Life*, 7(3), 2001.
- [140] A.C. Redfield. The biological control of chemical factors in the environment. *American Scientist*, 46(3):205–221, 1958.
- [141] R.P. Reid, P.T. Visccher, A.W. Decho, J.F. Stolz, B.M. Bebout, C. Dupraz, I.G. Macintyre, H.W. Paerl, J.L. Pinckney, L. Prufer-Bebout, T.F. Steppe, and D.J. DesMarais. The role of microbes in accretion, lamination, and early lithification of marine stromatolites. *Nature*, 406:989–992, 2000.
- [142] D. Robertson and J. Robinson. Darwinian Daisyworld. *Journal of Theoretical Biology*, 195:129–134, 1998.
- [143] R.B. Root. The niche exploitation pattern of the blue-gray gnatcatcher. *Ecological Monographs*, 37:317–350, 1967.
- [144] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D.E. Rumelhart and J.L. McClelland, editors, *Parallel Distributed Processing*, pages 318–362. 1986.
- [145] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.
- [146] D.E. Rumelhart and J.L. McClelland. *Parallel Distributed Processing*. MIT Press, Cambridge, MA, 1986.
- [147] L.C. Rutherford, A. DeWan, H. Lauer, and G.G. Turrigiano. BDNF has opposite effects on the quantal amplitude of pyramidal neuron and interneuron excitatory synapses. *Neuron*, 21(3):521–530, 1998.

- [148] P.T. Saunders. Evolution without natural selection: further implications of the Daisyworld parable. *Journal of Theoretical Biology*, 166:365–373, 1994.
- [149] S. H. Schneider, editor. *Climatic Change Volume 52 Number 4*. Springer, Netherlands, 2002.
- [150] S. H. Schneider, editor. *Climatic Change Volume 58 Number 1-2*. Springer, Netherlands, 2003.
- [151] S.H. Schneider and P. J. Boston, editors. *Scientists on Gaia: Papers delivered at the American Geophysical Union's annual Chapman Conference in March, 1988*. MIT Press, Cambridge, MA:, 1991.
- [152] D.W. Schwartzman and T. Volk. Biotic enhancement of weathering and surface temperatures on Earth since the origin of life. *Palaeogeogr. Palaeoclimatol. Palaeoecol.*, 90:357–371, 1991.
- [153] D.W. Schwartzman and T. Volk. Biotic enhancement of weathering and the habitability of Earth. *Nature*, 340:457–460, 1989.
- [154] J.R. Searle. Minds, brains and programs. *Behavioral and Brain Sciences*, 3:417–424, 1980.
- [155] A. Slocum, D. Downey, and R. Beer. Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 6: Proceedings of the 6th Intl. Conf. on Simulation of Adaptive Behavior*, Cambridge, MA, 2000. MIT Press.
- [156] E. Sober and D. S. Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Harvard University Press, Cambridge, MA, 1997.
- [157] M. Staley. Darwinian selection leads to Gaia. *Journal of Theoretical Biology*, 218:35–46, 2002.
- [158] M. Stemmler and C. Koch. How voltage-dependent conductances can adapt to maximise the information encoded by neuronal firing rate. *Nature Neuroscience*, 2(6):521–527, 1999.
- [159] S. Stöcker. Regarding mutations in Daisyworld models. *Journal of Theoretical Biology*, 175:495–501, 1995.

- [160] S.H. Strogatz. *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering*. Perseus Books, Cambridge, MA, 1994.
- [161] T. Sugimoto. Darwinian evolution does not rule out Gaia hypothesis. *Journal of Theoretical Biology*, 218:447–455, 2002.
- [162] W. Swenson, J. Arendt, and D.S. Wilson. Artificial selection of microbial ecosystems for 3-chloroaniline biodegradation. *Environmental Microbiology*, 2(5):564–571, 2000.
- [163] W. Swenson, D.S. Wilson, and R. Elias. Artificial ecosystem selection. *PNAS*, 97:9110–9114, 2000.
- [164] E. Thelen and L.B. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press, Cambridge, MA, 1996.
- [165] A. Thompson. Temperature in natural and artificial systems. In P. Husbands and I. Harvey, editors, *Proceedings of the 4th European Conference on Artificial Life (ECAL'97)*, pages 388–397. MIT Press, 1997.
- [166] P. Trojan. *Ecosystem Homeostasis*. Dr W. Junk, The Hague, 1984.
- [167] G.G. Turrigiano. Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends in Neuroscience*, 22:221–228, 1999.
- [168] G.G. Turrigiano, L.F. Abbott, and E. Marder. Activity-dependent changes in the intrinsic properties of cultured neurons. *Science*, 264:974–977, 1994.
- [169] G.G. Turrigiano, K.R. Leslie, N.S. Desai, L.C. Rutherford, and S.B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical pyramidal neurons. *Nature*, 391:892–895, 1998.
- [170] G.G. Turrigiano, G. Le Masson, and E. Marder. Selective regulation of current densities underlies spontaneous changes in the activity of cultured neurons. *J. Neurosci.*, 15:3640–3652, 1995.
- [171] G.G. Turrigiano and S.B. Nelson. Hebb and homeostasis in neuronal plasticity. *Current Opinion in Neurobiology*, 10:358–364, 2000.
- [172] G.G. Turrigiano and S.B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5:97–107, 2004.

- [173] T.J. van Gelder. What might cognition be, if not computation? *Journal of Philosophy*, 91:345–381, 1995.
- [174] V.I. Vernadsky. The biosphere and the noosphere. *American Scientist*, 33(1):1–12, 1945.
- [175] V.I. Vernadsky. *The Biosphere*. Copernicus, New York, 1998. English translation, Russian published 1926.
- [176] T. Volk. *Gaia's Body: Toward a Physiology of Earth*. Springer-Verlag, New York, 1998.
- [177] T. Volk. Toward a future for Gaia theory. *Climatic Change*, 52:423–430, 2002.
- [178] T. Volk. Natural selection, Gaia, and inadvertent by-products. *Climatic Change*, 58:13–19, 2003.
- [179] V. Volterra. Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. *Mem. R. Accad. Naz. dei Lincei. Ser. VI*, 2, 1926.
- [180] W. von Bloh, A. Block, and H.J. Schellnhuber. Self-stabilization of the biosphere under global change: a tutorial geophysiological approach. *Tellus*, 49B:249–262, 1997.
- [181] M.J. Wade. An experimental study of group selection. *Evolution*, 31:134–153, 1977.
- [182] A.J. Watson. Coevolution of the Earth's environment and life: Goldilocks, Gaia, and the anthropic principle. In G.Y. Craig and J.H. Hull, editors, *James Hutton - Present and Future*, pages 75–88, London, 1999. Geological Society.
- [183] A.J. Watson and J.E. Lovelock. Biological homeostasis of the global environment: the parable of Daisyworld. *Tellus*, 35B:284–289, 1983.
- [184] S.L. Weber. On homeostasis in Daisyworld. *Climatic Change*, 48(2-3):465–485, 2000.
- [185] S.L. Weber and J.M. Robinson. Daisyworld homeostasis and the Earth system. In J. Miller, P. Boston, S. Schneider, and E. Crist, editors, *Scientists Debate Gaia: The Next Century*, Cambridge, MA, 2004. MIT Press.
- [186] G.C. Williams. *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton University Press, Princeton, NJ, 1966.

- [187] G.R. Williams. *The Molecular Basis of Gaia*. Columbia University Press, New York, 1996.
- [188] D.H. Wolpert and W.G. Macready. No free lunch theorems for search. Technical Report SFI-TR-95-02-010, Santa Fe Institute, 1995.
- [189] D.H. Wolpert and W.G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(67), 1997.
- [190] A.J. Wood, G.J. Ackland, and T.M. Lenton. Mutation of albedo and growth response produces oscillations in a spatial “Daisyworld”. Submitted to *Journal of Theoretical Biology*.
- [191] J.P. Wright, C.G. Jones, and A.S. Flecker. An ecosystem engineer, the beaver, increases species richness at the landscape scale. *Oecologia*, 132:96–101, 2002.
- [192] B. Yamauchi and R.D. Beer. Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3):219–246, 1994.
- [193] X. Zeng, R. Pielke, and R. Eykholt. Chaos in Daisyworld. *Tellus*, 42B:309–318, 1990.
- [194] W. Zhang and D.J. Linden. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews Neuroscience*, 4:885–900, 2003.