

Computer Models of The Evolution of Language
and Languages

DANIEL JACK LIVINGSTONE

A thesis submitted in partial fulfilment of the
requirements of the University of Paisley for
the degree of Doctor of Philosophy

September 2003

Abstract

The emergence and evolution of human language has been the focus of increasing amounts of research activity in recent years. This increasing interest has been coincident with the increased use of computer simulation, particularly using one or more of the methods and techniques of ‘Artificial Life’, to investigate a wide range of evolutionary problems and questions. There is now a significant body of work that uses such computer simulations to investigate the evolution of language.

In this thesis a broad review of work on the evolution of language is presented, showing that language evolution occurs as two distinct evolutionary processes. The ability to use language is clearly the result of biological evolution. But the changes that occur over time to all spoken languages can also be viewed as being part of a process of cultural evolution. In this thesis, work using artificial life models to investigate each of these processes is reviewed. A review of the methods and techniques used in artificial life is also presented early in the work.

A novel model is developed which is used to explore the conditions necessary for the evolution of language. Interesting results from initial tests of the model highlight the role of redundancy in language. From these initial tests, the model is further developed to explore the biological evolution of the human capacity for language. One significant outcome of this work is to highlight the limitations of the model for developing, and especially for ‘proving’, particular theories on how or why *Homo sapiens* alone evolved language. This is tied to a brief review showing that this weakness is not one specific to this particular model, but may be one that is possessed by all artificial life models that try to explain the origins of language.

With further minor modifications to the model, the focus is shifted to the evolution of languages and language diversity. In comparison with some of the earlier conclusions, this work emphasises the positive contribution to ongoing scientific debate that is possible using computer simulations. In this case, experiments using the model focus on whether social and/or linguistic benefits are required in explanations of language change. A review and debate is then presented on work that contradicts our findings. Further corroboration of our conclusions is then gained by conducting a similar experiment using a different computer model.

The key contributions of this interdisciplinary work are: first, in detailing some of the unique problems and issues inherent in using computer models specifically for modelling the evolution of language; second, in emphasising the importance of redundancy in language evolution; and finally, in adding to the current debate on whether the evolution of languages can be viewed as a form of adaptively neutral evolution.

Dedication

For dad.

Acknowledgements

That this thesis is now finally complete is thanks, in great part, to the help, comments, advice, and support provided by many people. I should like to take this opportunity to thank some of them.

First, I must acknowledge the contribution of Professor Colin Fyfe. Without his many words of encouragement, regular cajoling and occasional inquisitions it is doubtful that I would ever have completed. Thank you Colin for your great patience in seeing this through. I should also like to thank the School of Computing, and its predecessor departments for paying my fees and providing financial support.

Being the sole researcher at an institution working on a particular topic can be lonely at times, so I would like to thank my departmental colleagues for making it considerably less so – particularly Darryl Charles, Stephen McGlinchy, Donald McDonald, Bobby Geary and Douglas Wang.

When I started out on my research I had little idea what I was getting myself into. A number of people I met at conferences and visits to other institutions helped me find out. To list them all would not be possible, but here I would like to offer thanks to Jason Noble, a friendly soul, through whom I met, and had a number of interesting chats with Inman Harvey, Ezequiel di Paolo and Seth Bullock. Simon Kirby is due thanks for numerous pointers, listening to my moaning, and those proceedings that I have yet to return. Mike Oliphant is thanked for giving me a second chance at explanation. Thanks are due too to those other friendly souls, Bart De Boer, Paul Vogt and Tim Taylor, and the members of the Edinburgh LEC group.

David Hales, Jeremy Goslin and Chris Douce are all due an apology – I really should have returned more emails. Sorry. Dave also pointed me to some interesting papers, and told me lots about Artificial Life long before I ever thought I'd end up working with it myself.

The final shape of this thesis is also thanks to the hard work of many anonymous reviewers whom I would also like to thank. The very un-anonymous Angelo Cangelosi and Jim Hurford gave this thesis its very final review and examination – and I should like to thank them both for their careful reading and positive comments.

Away from the world of academia, thanks are due to those great imbibers, with whom many an enlightening conversation was had: Arthur McGready and Alasdair Duncan. Thanks also to the Players of Games and the Writers of Books: Gary Devine, Dougie Gaylor, Tom McGlaughlin, Pete McMinn, Phil Raines and Gary Gibson.

And finally, I am eternally indebted to Bronwen and Morgan for their love and support. You make it all worthwhile. Thank you.

Contents

Chapter 1 Introduction.....	1
1.1 <i>Outline.....</i>	1
1.2 <i>Artificial Life and Artificial Societies</i>	2
1.3 <i>A Note on Previous Publications</i>	3
Chapter 2 The Evolution of Communication and Language.....	5
2.1 <i>Introduction.....</i>	5
2.2 <i>Language and Communication</i>	6
2.2.1 <i>Features of Language</i>	8
2.3 <i>The Evolution of Communication</i>	10
2.4 <i>The Biological Evolution of Language</i>	13
2.4.1 <i>Evidence for the Innateness of Language</i>	13
2.4.2 <i>Evolution of Linguistic Ability.....</i>	14
2.4.3 <i>Questioning the Innate Linguistic Ability.....</i>	16
2.4.4 <i>Why and How? Just So Stories and Grand Theories</i>	18
2.5 <i>The Cultural Evolution of Languages.....</i>	21
2.5.1 <i>Language Change.....</i>	21
2.5.2 <i>Explanations of Language Change</i>	23
2.5.3 <i>Language Change and Evolutionary Theory</i>	25
2.5.4 <i>Cultural Evolution.....</i>	26
2.5.5 <i>Language Ecology and Species</i>	27
2.5.6 <i>Neutral Evolution.....</i>	29
2.5.7 <i>Co-evolution of Languages and Man.....</i>	30
2.6 <i>Summary</i>	30
Chapter 3 Artificial Life: Computational Modelling and a Methodological Approach	31
3.1 <i>The Artificial Life Approach</i>	31
3.1.1 <i>Agents</i>	31
3.1.2 <i>Emergence.....</i>	32
3.1.3 <i>Life-Like</i>	32
3.2 <i>Methodologies for Applying Artificial Life.....</i>	33
3.2.1 <i>Science without method?</i>	34
3.3 <i>Improved Methodologies and Practices</i>	35
3.3.1 <i>A Methodology Emerges</i>	35
3.3.2 <i>Not Just Biology</i>	37
3.3.3 <i>ALife for Novel Models and Theories.....</i>	37
3.4 <i>ALife in Practice</i>	39
3.4.1 <i>ALife Research vs ‘Real’ Research</i>	39
3.4.2 <i>Principles of ALife Model Building</i>	40

3.4.3 Adapting Methods.....	41
3.4.4 Implementing an ALife model.....	42
3.5 <i>Experimentation</i>	43
3.5.1 Validation.....	43
3.5.2 Verification.....	44
3.5.3 The (Un-)Importance of Surprise.....	46
3.6 <i>Explanation</i>	47
3.6.1 Model and Theory.....	48
3.6.2 Validation of Results.....	48
3.7 <i>Limitations Of Artificial Life</i>	49
3.7.1 ALife vs. Mathematical Models.....	49
3.7.2 ALife Models as Proof.....	50
3.7.3 ALife Models as Evidence.....	51
3.7.4 ALife and Quantitative Results.....	52
3.8 <i>ALife and the Evolution of Language</i>	53
3.8.1 ALife and the Biological Evolution of Language.....	53
3.8.2 ALife and the Cultural Evolution of Languages.....	56
3.8.3 Coevolution of Language Ability and Languages.....	57
3.9 <i>Summary</i>	57
Chapter 4 An Artificial Life Model for Investigating the Evolution of Language	59
4.1 <i>Introduction</i>	59
4.2 <i>Computational Models of Language Learning Populations</i>	59
4.2.1 Evolution of Language Ability.....	61
4.3 <i>Modelling Human Language</i>	62
4.4 <i>An Artificial Neural-Network Based Language Agent</i>	64
4.5 <i>Evaluating Communication in Artificial Populations</i>	65
4.6 <i>Modelling Language-Physiology Coevolution</i>	66
4.7 <i>A Language Agent</i>	66
4.7.1 Signal Production.....	68
4.7.2 Signal Interpretation.....	68
4.7.3 Learning.....	69
4.7.4 A Note on Terminology.....	69
4.7.5 Representational Capacity.....	69
4.7.6 A Bias for Language Learning.....	70
4.7.7 The Individual Language Agent – 1.....	71
4.7.8 The Individual Language Agent – 2.....	71
4.8 <i>Experiment 1: Language Negotiation</i>	72
4.8.1 Emergence of a Language Bias.....	72
4.8.2 Success of Language Negotiation.....	73
4.8.3 Interpreting the Results.....	75
4.9 <i>Experiment 2: Spatially Arranged Populations</i>	75
4.9.1 Emergence of Dialects.....	77

4.10 Conclusion	78
Chapter 5 Modelling the Biological Evolution of Language.....	80
5.1 <i>The Evolution and Emergence of Language</i>	80
5.2 <i>Modelling the Biological Evolution of Language Ability</i>	81
5.2.1 Heterogeneous Language Abilities	81
5.2.2 Comprehension Leads Production	83
5.2.3 Population Generations and Replacement	84
5.2.3.1 Genetic Representation and Reproduction.....	84
5.2.3.2 Selection, Mating and Replacement	86
5.2.3.3 Crossover and Mutation.....	87
5.3 <i>Modelling Language-Physiology Coevolution, Part I</i>	88
5.3.1 Crossover and Mutation	88
5.3.2 Parameters and Settings	89
5.3.3 Discontinuous Evolution of Signalling Ability.....	89
5.4 <i>Modelling Language-Physiology Coevolution, Part II</i>	92
5.4.1 Crossover and Mutation	92
5.4.2 Parameters and Settings	92
5.4.3 Continuous Evolution of Signalling Ability	93
5.5 <i>The Effect of Neighbourhood Size</i>	94
5.5.1 Parameters and Settings	94
5.5.2 The Effect of Large Neighbourhoods	95
5.6 <i>Coevolution with a Costly Language Ability</i>	96
5.6.1 Parameters and Settings	96
5.6.2 Spatial Selection and Costly Language Ability	96
5.7 <i>Discussion</i>	101
5.7.1 Continuous versus Discontinuous Evolution of Language.....	101
5.7.2 Spatial Selection.....	101
5.7.3 Redundancy and Linguistic Ability	102
5.7.4 Investigating The Adaptive Benefits of Language	103
5.7.5 Embodied Communication	104
5.7.6 Limitations and Shortcomings	107
5.8 <i>Summary</i>	108
Chapter 6 The Cultural Evolution of Language and The Emergence and Maintenance of Linguistic Diversity	111
6.1 <i>Introduction</i>	111
6.2 <i>An Initial Test</i>	112
6.3 <i>Human Linguistic Diversity</i>	113
6.3.1 Patterns of Diversity	113
6.3.2 Linguistic Boundaries	114
6.4 <i>Analytical Models of Linguistic Diversity</i>	115
6.4.1 The Niyogi-Berwick Model.....	116
6.4.2 Criticism of the Niyogi-Berwick Model.....	117

6.4.3 The Cavalli-Sforza and Feldman Models	119
6.4.4 Summary of Analytical Models	120
6.5 <i>Modifying the Artificial Life Model for Exploring the Emergence of Linguistic Diversity</i>	121
6.6 <i>Experiment 1: Emergence and Maintenance of Dialects</i>	122
6.6.1 Experimental Setup	122
6.6.2 Results	123
6.6.2.1 Visualisation	123
6.6.2.2 Experimental Results	124
6.6.2.3 Measuring Diversity	125
6.6.3 The Effect of Neighbourhood Size	126
6.6.4 Analysis	127
6.6.4.1 Maintaining Diversity	127
6.6.4.2 Human linguistic diversity: A comparison	129
6.7 <i>Experiment 2: Diversity from Homogeneity</i>	131
6.7.1 Experimental Setup	132
6.7.1.1 Noisy Learning	132
6.7.1.2 Innate Bias	133
6.7.2 Results	133
6.7.3 Discussion	133
6.8 <i>Summary</i>	133

Chapter 7 Cultural Evolution in an Agent Based Model of Emergent Phonology 135

7.1 <i>Artificial Life and Micro-Simulation Models of Linguistic Diversity</i>	135
7.1.1 Functional Requirements for Diversity	135
7.1.2 Other Models of the Evolution of Linguistic Diversity	137
7.1.3 Related Models	137
7.2 <i>Dialect in an Agent Based Model of Emergent Phonology</i>	138
7.2.1 de Boer's Model of Emergent Phonology	139
7.2.2 Experimental Setup	142
7.2.3 Results	143
7.3 <i>Discussion: Towards a Modified-Neutral Theory of Language Change and Diversity</i>	145
7.3.1 Neutral Evolution Revisited	145
7.3.2 Adaptive, Maladaptive and Neutral Change	146
7.3.3 Universality and Uniformitarianism	149
7.3.3.1 Universality	149
7.3.3.2 Uniformitarianism	149
7.3.4 Relativity and Non-Uniformitarianism	150
7.3.4.1 Linguistic Relativity Hypothesis	150
7.3.4.2 Non-Uniformitarianism	150
7.3.4.3 Directed Change in Sound Systems	151
7.3.5 Neutral Change in a Relativistic, Non-Uniformitarian World	151
7.3.5.1 Neutral Change In Sound Systems	151
7.3.5.2 Neutral Networks in Language Evolution	152
7.3.5.3 Neutral and Adaptive Evolution	153

7.4 <i>A Modified-Neutral Theory of Linguistic Evolution</i>	153
7.5 <i>Conclusions</i>	154
Chapter 8 Conclusions	155
8.1 <i>Artificial Life and the Evolution of Language</i>	155
8.2 <i>Methodological Approaches</i>	156
8.3 <i>Redundancy in Language Evolution</i>	157
8.4 <i>Language Change and Neutral Evolution</i>	158
8.5 <i>Future Directions</i>	159
Appendix A	160
Appendix B Colour Plates	163
References	171

Chapter 1 Introduction

1.1 Outline

In this work I apply computational modelling methods to study some questions on the evolution of language. The computational approach taken is an ‘Artificial Life’ one, and just what this entails is necessarily reviewed early in this work, in Chapter 3, and is also outlined in brief below.

It is also necessary to review some of the vast array of work aimed at answering the questions, big and small, that exist about the evolution of language: why did the ability of humans to use language evolve, and why only for *Homo sapiens* and how did the capacity for language evolve? Other, subtler and more fundamental, questions emerge: what *is* language? what is actually innate, and biologically evolved, in the human capacity for language and what is simply the result of cultural learning processes over many generations?

Thus, two distinct forms of language evolution become apparent – the biological evolution of some innate linguistic ability, and the cultural evolution of specific languages and language families. After the terminology of Hurford (1999), who distinguishes these as the evolution of language and the evolution of languages, these two forms are here termed the Evolution of Language (EoL) and the evolution of languages (eol) respectively. (By convention, in the remainder of this thesis these two types of evolution of language will be referred to by their abbreviated forms. The long hand phrase ‘evolution of language’ will be used when referring to the evolution more generally.) A broad picture of current thought on these two forms of linguistic evolution is given in the literature review (Chapter 2).

In Chapter 4, I describe the basic details of the computational model I will use in my own investigations. Some tests are carried out on the model to demonstrate its workings. Then, in Chapter 5, the model is used to demonstrate the continuous evolution of linguistic ability within a population of language users. I will argue that this demonstration is relevant to the continuity-discontinuity debate that exists on the EoL, and emphasises some of the conditions necessary for the EoL.

In contrast, in Chapter 6, the model is adapted to investigate the eol. Additionally, a model of emergent phonology is used to support the results gained from the existing computational model. As with EoL, there is current debate in studies of the eol on

how computational models may aid in the formulation of arguments and in the provision of additional evidence. The results of the experiments are analysed and exactly what they say about the evolution of language explored in the final chapter of the thesis.

An important goal of this research is to explore the suitability of Artificial Life techniques for study of the evolution of language and languages, and accordingly a critical eye is cast over the work, with this in mind, in the conclusion. While the number of investigators now using Artificial Life methods to investigate the evolution of language seems to indicate that Artificial Life is a useful approach, some shortcomings in my own work and a few in the other works will be highlighted.

As indicated above, a quick introduction to Artificial Life is in order. This is given below, and will be reviewed in more detail in Chapter 3.

1.2 Artificial Life and Artificial Societies

Artificial Life, hereafter ALife, is a comparatively recent offshoot of Artificial Intelligence, with its roots in the unrelated fields of distributed Artificial Intelligence (DAI) and computer simulation, and biology. The term Artificial Life is itself quite new, coined by Chris Langton, organiser of the first workshop on Artificial Life (Langton, 1989). The workshop was organised to bring some unity to the fragmented literature which existed on biological modelling and simulation. Through the workshop, Langton reports that some consensus of the “essence” of Artificial Life emerged:

“Artificial Life involves the *realization* of life-like behaviour on the part of man-made systems consisting of *populations* of semi-autonomous entities whose *local interactions* with one another are governed by a set of *simple rules*... high-level dynamics and structures observed are *emergent* properties, which develop over time from out of all of the local interactions...”

(Langton, 1989, p. xxii, author’s emphasis)

This states that Artificial Life is simply an advanced modelling technique – one where the behaviour of whole populations can be observed, although the rules for how the population might behave are not explicitly detailed within the model. Instead, the behaviour of individuals within the population are detailed, this hopefully

being less reliant on possibly incorrect and simplistic assumptions. Through computational processes, individuals interact repeatedly and over time and the results at the population level are observed.

Although seeming new, the computational approach to micro-modelling had in fact been explored before by scientists in other communities, with some early ground-breaking work in social modelling in particular (Schelling, 1978). This idea has been remarkably successful, and ALife techniques are now applied to a wide range of problems beyond biological and social modelling, and a number of international conferences are held annually with ALife as a central topic.

Societies and journals exist to support researchers using these new computational techniques in various fields – the *Artificial Life* journal, *The Journal of Artificial Societies and Social Simulation*, the International Society of Adaptive Behaviour, and more. Also, as evidenced by this thesis, and a number that have gone before it, ALife is being increasingly applied to the complex issues surrounding the evolution of language (for further examples, also see (Hurford et al., 1998; Dessalles and Ghadapour, 2000; Cangelosi and Parisi, 2002)).

1.3 A Note on Previous Publications

Much of the work presented in this thesis has been previously published in different form.

The material of Chapters 2 and 3 is drawn from the literature review, which was presented as part of the transfer report on the 28th of January 1999. Chapter 3 also includes much of the article (Livingstone, 1999a), as published in the C.I.S. departmental journal at the University of Paisley. This has been substantially altered since, with a number of additional sections.

The description of the model presented in Chapter 4 is drawn from, again with additional material, previously published descriptions in (Fyfe and Livingstone, 1997), (Livingstone and Fyfe, 1998a; Livingstone and Fyfe, 1998b; Livingstone and Fyfe, 2000). The experiments of Chapter 5 are based on the same works, but with considerably more detail and expansion.

The bulk of the material of Chapter 6 was previously presented in different forms in (Livingstone, 1999b; Livingstone and Fyfe, 1999a; Livingstone and Fyfe, 1999b; Livingstone, 2000b) and (Livingstone, 2002), from initial exploratory experiments to later developments. As a matter of course, many of these reproduce selected details

of the model construction, as presented here in Chapter 4. Some of the material detailing work with an alternative phonological model, as presented in Chapter 6, was previously presented in (Livingstone, 2000a).

Chapter 2 The Evolution of Communication and Language

2.1 Introduction

The aim of this chapter is to provide a grounding in the domain of study – that of the evolution of language. Some definitions of terms that will be used throughout the thesis will be given, and some of the many still open questions about the evolution of language will be identified. In particular, the problems which this thesis will later investigate will be detailed.

The evolution of language and communication is an inherently inter-disciplinary field of research, of interest to researchers in anthropology, archaeology, biology, linguistics, psychology, physiology and neuro-physiology (e.g. Bickerton, 1984; Pinker, 1994; Dunbar, 1996; Mithen, 1996; Deacon, 1997; Hurford et al., 1998).

Thus an area of research has been staked out that, while apparently focussed, inhabits an incredibly diverse space. In this review, I draw on sources from many different fields, looking first in broad terms at language and communication – specifically at what differentiates human language from other systems of communication.

Then I review some theoretical work on the evolution of communication. This concentrates on animal communication rather than human language, but provides a base for much work on the EoL. Here we see links to another fundamental problem, one which has also been the subject of much ALife investigation – the evolution of cooperation (Axelrod, 1984).

I then turn my attention to work on the biological evolution of language. First I introduce some of the evidence for a biological basis for language, and some of the major theories on innate linguistic abilities. I proceed to review some work questioning the precise nature of the innate language ability before looking at some attempts to explain why and how language evolved.

Finally, I move from work on the biological EoL to look at the cultural eol. Here the concern is with the historical development of particular languages, how languages change over time, rather than how and why humans have gained the ability to use language. However, there is some degree of crossover between these questions. As I will show when questioning the degree of innateness of language, some recent work

on culturally evolving linguistic systems provides a challenge to current ideas about the human biologically innate language ability.

This chapter provides the background knowledge required before meaningful, and hopefully useful, ALife experiments can be carried out.

As a last word in this introduction, note that, in this chapter, I will introduce a number of terms with special meaning within linguistics. This is unavoidable – linguists use a great deal of jargon, and have created a rich language for discussing language itself. I have deliberately kept the use of such jargon to a minimum, emphasising only terms of special importance or those which will reappear later in this thesis.

2.2 Language and Communication

A spectacular variety of means of communication between animals (of the same or different species) have evolved, human language being but one, albeit unique in a number of respects. Animal signalling can take many forms and signals can use sound, vision, touch and even scent or taste as a means of signal transmission/reception. Before we can proceed to model the evolution of language, it is required that some understanding of the distinction between non-linguistic communication and language is gained. Then we can hope that the models which are developed are relevant to the evolution of language. This requirement for making a distinction is expanded upon in the next chapter.

In this work I side-step a rather fundamental, and surprisingly intractable, problem for theoretical biologists – that of defining what communication actually is. To get some idea of this problem consider first a simple view of communication, commonly assumed:

A signal can be said to be communicative where the originator of the signal intentionally uses it to transmit information to the receiver.

There are significant problems with such a simple definition. First, if signalling is viewed simply as the transfer of (mis)information from one individual to another, then signals may be given voluntarily or involuntarily (Krebs and Dawkins, 1984). Bullfrogs croak to advertise their quality as mates. Can we say that any bullfrog intends to advertise itself as being of low quality? Instead the attempt of the female to

determine which will be a good mate by listening to the croaks is what Krebs and Dawkins term ‘mind-reading’. The mind-reading perspective views signals as an attempt by a signal receiver to determine some information about some target which is emitting the signal – but the sender may not intend for the signal to be observed, or for a particular inference to be made from it.

Any attempt to cite purpose on behalf of the receiver to circumvent this problem is similarly limited – a complimentary ‘manipulation’ perspective views signals as an attempt to exploit another individual on behalf of the signaller.

Attempts to define communication in terms of information transfer (e.g. Grice, 1969) have problems too. For example, how should deceptive signals be classified? Additionally, in a theoretical study, Di Paolo (1999b) notes that communication can be useful even where the signals have no informational content. (In linguistics such non-informational communication may be termed *phatic* where the intent of the communication is, e.g., to establish a common ‘mood’).

Thus neither intent nor informational content can be assumed when attempting to define communication. Another approach, again from biology, developed by Millikan, is to declare signalling systems to exist only where selection has adapted the communicating agents to play their *cooperative* roles in signalling episodes (Bullock, 2000). This solves the problems noted above, but leaves many apparent examples of communicative behaviour outside the definition of signalling systems. Bullock further notes that in many cases in biology, communication involves situations in which the interests of the agents coincide and conflict in parts. Further discussion of many of these issues can be found in theses and papers by Bullock and Noble (Bullock, 1997; Bullock, 2000; Noble, 2000).

Yet, however imperfect, limiting the definition of communication to intentional signalling which aims to benefit both sender and receiver of a signal – limiting it to deliberate, cooperative signalling – can provide a useful start point for studying the evolution of communication. That a single, universal, definition of signalling has yet to be found that satisfies the whole breadth of studies into communication necessitates our selection of a more limited definition. This simplifying definition is assumed as a given in my later experiments on the evolution of language. Relating to this, however, there is more discussion on the functional benefits of language, and other issues, in the remainder of this and in the next chapter.

However, my goal here is to distinguish between linguistic and non-linguistic communication. Rather than struggle with definitions, one way of doing this is to compare the characteristic features of language versus those of other communication systems.

2.2.1 Features of Language

Charles Hockett (1960) suggested that language evolved gradually over time, suggesting a continuity with non-human communication systems. Hockett broke language down into thirteen components.

It is worthwhile listing these features, representing one of the first modern attempts to identify precisely what is unique about human language compared with other forms of communication. The features are chosen as *characteristic* of human language – exceptions may exist, and some of these will be mentioned.

The features of language which Hockett identifies are:

1. Vocal-auditory channel: sounds are used for language transmission.
2. Broadcast transmission and directional reception: signals can be heard by anyone within range, and the signals can be localised.
3. Rapid-fading: speech is transitory in nature, signals exist for only a brief period of time.
4. Interchangeability: language users can produce any message that they are capable of understanding.
5. Total feedback: speakers can hear and are able to reflect upon everything they say.
6. Specialization: the sound waves' only function is to signal meaning.
7. Semanticity: elements of a signal convey meaning through stable signal-meaning associations.
8. Arbitrariness: the signal-meaning associations are arbitrary.
9. Discreteness: language is based upon the use of a small set of discrete sound elements.
10. Displacement: it is possible to use language to refer to events displaced in both time and space from the speaker.
11. Productivity: it is possible to produce (and understand) an infinite range of meanings, using existing elements to produce new sentences.

12. Traditional transmission: language is transmitted by a cultural learning process rather than genetic transmission of innate calls.
13. Duality of patterning: the sounds used in language do not convey meaning, but combined in different ways to form words they do.

Of these, the first three features relate to the auditory nature of language. Yet it has been observed that communities of deaf children will form their own natural sign-language (Kegl and Iwata, 1989). This challenges the significance of, at least, the first two features.

While many of these features can be found in various animal communication systems, only human language contains all thirteen of these features. Hauser (1996) points out that no insight is provided into the functional significance of these features, and that they are not weighted at all by their relative importance. So we do not know which are the *significant* features of language simply by scanning this list.

It was considered that only human language featured traditional transmission or duality of patterning, but this is now known not to be the case. However, like displacement, which other than human language is only known to exist in the dance ‘language’ of bees, these features are very rare in natural communication systems. And while the thirteen features are characteristic of human language, they are not definitive – sign-language being an obvious exception.

In building an ALife model to help investigate the EoL, it should not be necessary to include all of these features, but unless the model is able to include some significant features it may not be a useful model for its intended purpose. To emphasise the features of greatest significance, it is helpful to compare the previous list with another from a more recent source. A textbook listing (from Sternberg, 1996) states that the characteristic features of language are that it is:

- Communicative: language is used to transfer and share information.
- Arbitrarily symbolic: the sounds and symbols used in language are primarily arbitrary, having no set relation to the underlying meanings.

The signals used in human language for communication bear no relation to the concepts being communicated, other than a purely arbitrary pairing of sound and concept. This principle of the ‘arbitrariness of the sign’ is expounded in the work of Ferdinand de Saussure (1916), and the idea can be traced back to Aristotle and beyond.

The term ‘Saussurean communication’ refers to communication systems where concepts and signs are related to one another by such arbitrary relationships.

- Regularly structured and structured at multiple levels: words are constructed from phonemes, phrases from words and sentences from phrases. Regular structures occur at these different levels.
- Generative and productive: it is possible to use language to express effectively infinite numbers of new meanings. Learned language can be used to produce original statements.

Structure exists in language at multiple levels. Morphology looks at the internal structure of words, syntax at how words are combined to form phrases and sentences. That language has hierarchical structures allows a finite number of language elements to be used in the generation of infinite numbers of sentences and meanings.

- Dynamic: language is not fixed, but develops and changes with time.

Language changes and develops over time – both for individual language users as their own language skills develop, and for communities as languages change over time. A considerable body of work exists studying the development of language in children (e.g. Elliot, 1981).

More recently, Carstairs-McCarthy (1999) highlights three key features of language. Beside the previously noted duality of patterning, two features are proposed which are not normally considered characteristic features of language. The first of these is the number of words with distinct meanings that may exist in any language. Here there is a significant difference between human language and animal communication systems. The other feature proposed is the distinction between particular syntactic categories, introducing different *types* of meanings – this last feature seemingly peculiar to the author’s theory of the evolution of language.

2.3 The Evolution of Communication

Amongst the vast array of animal communication systems are such diverse systems as bee dances to convey information about the location of pollen-rich flowers; ant pheromone communication to lead other ants from the same colony to food and back to the nest; and vervet monkey alarm calls used to warn other vervets of the presence of specific types of predators – eagle, leopard or snake.

Biologists have been concerned for some time with how such systems could evolve.

Lorenz (1966) and Tinbergen (1976) claim that many signals may have evolved from incidental movements or responses of actors which happen to pass information to reactors. What was once an accidental pairing of some action with a response is repeated and reinforced over time until it becomes a conventionalised communicative signal and response – a process termed ritualization. This can succeed as selection favours those able to interpret such signals, e.g. to anticipate (and avoid) attack, and those who produce the signals, e.g. to scare off others with threat of attack and hence use less energy than in an actual attack. How a signal is ritualised may vary with the purpose of the signal and the rewards for actor and reactor. Ritualization may also occur to reduce ambiguity in signal reception, increasing the distinctiveness of the signalling movement.

In some cases communication serves a clearly co-operative purpose – vervets warning each other of predators, or bees indicating where pollen may be found – and studies in the evolution of communication often focus on communication that is co-operative in nature (Hauser and Marler, 1999). Assuming that communication is cooperative leads naturally on to another problem – that of the evolution of cooperation. If communication is cooperative, then it might be expected to emerge in similar circumstances to those that can support the evolution of cooperation.

The fundamental problem for the evolution of cooperation is that in many cases cooperation does not appear to be an evolutionary stable strategy. If cooperation emerges in a population, any defectors that exist will be able to exploit the cooperation without rewarding those they exploit. Thus defection rather than cooperation will succeed in the population.

A number of mechanisms have been proposed to overcome this problem. Kin selection (Hamilton, 1964) is one – where cooperation occurs between kin, individuals within groups of kin which have evolved to cooperate will be fitter than those which live in non-cooperating family groups. Over time this can lead to a population where cooperative behaviour is the norm.

Zahavi (1979) proposed a handicap principle for honest signalling in sexual selection. This presumes that reactors will try to avoid manipulation by attending to honest indicators of qualities of size, strength, desire to fight, etc., ignoring other signals which then fall into disuse. Thus, surviving signals will be honest and costly to produce (thus only those individuals of higher quality will be able to produce higher quality signals). There should also be a direct link between signal design and quality

being signalled. For example the pitch calls of frogs, used to attract mates, are linked to body size.

Other mechanisms that can support the evolution of cooperation include spatial selection and the ability of individuals to remember which of the others is cooperative, and to act differentially based on past experience. These can both play a significant role in the iterated prisoner's dilemma (IPD) (Axelrod, 1984) where the individuals play not just once, but many times. Axelrod formalised a successful strategy for cooperative behaviour in the IPD – 'tit-for-tat' – where initially cooperative individuals either cooperate or defect depending on whether their partner previously cooperated or defected.

Spatial selection works as an alternative to memory, and can similarly favour the evolution of cooperation. Working with the assumption that individuals are limited in the amount of mobility they have, models have been built in which agents are placed in some spatial arrangement and where they are limited to interacting only with other agents nearby. In such circumstances, cooperation can succeed (Kirchkamp, 1996). Because of the spatial constraint, agents tend to repeatedly interact with the same few agents and this allows clusters of cooperation to form.

Some researchers are investigating how the evolution of non-cooperative communication (such as predator-prey signalling) can be understood and explained (e.g. Bullock, 2000). The concern here is that not all animal signalling is cooperative, and this is certainly something which can extend to human language.

A thorough and comprehensive review of the evolution of animal communication (including some material on the evolution of human language) can be found in Hauser (1996). Many other relevant topics are reviewed, including such aspects as how different environmental conditions shape the particular form of communication used – such as how the different acoustic properties of jungles versus those of open savannahs will affect the type and use of acoustic signals for communication.

This section has presented a brief review of two large topics of study – the evolution of communication and the evolution of cooperation. It is, I hope, enough to give the required base for the remainder of this thesis.

2.4 The Biological Evolution of Language

2.4.1 Evidence for the Innateness of Language

It is an unavoidable fact that language is in some way an innate, and hence biologically determined, human trait as emphasised by the following evidence.

Language is confined to only one species, amongst whom it is a universal trait, shared by all with only extraordinary cases providing exceptions. No monkey or ape, however well trained, is capable of acquiring language as well as a human child of around five years of age (although the precise capacity that apes *do* have for language is a current topic of debate (Savage-Rumbaugh, 2000)).

While the idea that primitive or ‘savage’ societies had equally primitive or savage languages was accepted in previous centuries (e.g. Rousseau, 1755; Burnett, 1787), this has proved to be a fallacy. Of all human societies and communities throughout the world, there is no culture lacking language, or possessing only a limited language that can truly be considered ‘primitive’ in its linguistic development (Pinker, 1994). Perhaps the closest examples to ‘primitive languages’ to be found are pidgin languages.

Pidgin contact languages are formed by migration of peoples with different languages into single communities, and a wide range of these have been studied including sign-language pidgins (Kegl and Iwata, 1989). Pidgins are grammatically weak, missing many aspects of grammar found in every other language. Yet within a small number of generations – at times only one – pidgins develop into *Creoles* (Bickerton, 1981). A creole is a descendant language of a pidgin, but one that is rich in grammar. This rapid development of grammar where there was none is a rich source of evidence on the innateness of grammar abilities in humans.

But what *is* the biological basis for language, and in what ways does it determine the shape and form of human languages?

The ability of children to quickly learn complex languages from fragmentary and incomplete evidence led Chomsky to argue that children are innately equipped with rules common to all languages - a Universal Grammar. While different languages have different grammars, the existence of some kind of Universal Grammar, UG, (Chomsky, 1972) is now widely accepted (e.g. Slobin, 1979; Pinker, 1994). The UG is believed to contain a number of rules and meta-rules for grammar. For example, phrase head and role-player ordering (e.g. verb-object, preposition-noun phrase,

adjective-complement) differs between languages, but remains constant in any single language. Thus, the UG contains a rule used in phrase construction, which allows either order but the same order throughout any language. As well as the UG, Chomsky proposed that there was an innate device for learning language.

2.4.2 Evolution of Linguistic Ability

The innate device for language learning was termed by Chomsky the Language Acquisition Device, LAD. The LAD has been the focus for much research, and despite the fact that a LAD would surely be an evolved solution for language learning, the evolution of the LAD and language itself is the subject of much debate in linguistics.

Chomsky himself rarely considers the problem of how the LAD may have evolved, and does not see the study of its evolution as being of value:

“it seems rather pointless... to speculate about the evolution of human language from simpler systems...”

Noam Chomsky, 1972, page 70

At times he appears to discount the idea that it evolved at all – rather he claims that it came about by chance, and only once it appeared fully formed did it become used for language.

Another view is offered by Bickerton (1990) who proposes that at some point in human evolution a basic form of language existed – one which he calls ‘protolanguage’. This protolanguage would have been lacking in a number of the features that characterise modern human language – principally it would have possessed a very limited grammar, consisting perhaps of only one or two-word utterances and incapable of complex sentence structure, or something akin to modern pidgins. A single ‘macro-mutation’ was then responsible for the evolution of fully developed language ability and the emergence of language.

Chomsky and Bickerton both present views of a sudden and discontinuous EoL. A number of other authors (e.g. Pinker and Bloom, 1992) argue instead that the EoL must have been gradual, and taken place over a longer time period, occurring in smaller steps. This is a continuous view of language evolution, and generally more in keeping with modern evolutionary theory.

Yet the idea that somehow language did not ‘evolve’ but it suddenly appeared fully-formed persists. Such arguments are often based on the principle of *exaptation*.

In evolutionary theory, adaptation is the process whereby organisms improve their fitness by gradual modification to better fit their environment. In contrast to this, Gould and Vrba (1982) coined the term exaptation to refer to the process of utilising some structure for some purpose other than that for which it originally evolved to serve. A good example of this might be penguins' use of wings for swimming.

Arguments against the idea that language is the end result of some sudden exaptation, rely on the number of costly adaptations that have occurred to support language. The larger brain necessary to enable speech costs more energy to maintain and requires a longer infancy to allow brain growth to complete (see, for example, Dunbar, 1996). The dropped larynx allows greater clarity and distinctiveness in speech, but may increase the risk of choking and was taken to be clear evidence of the evolution of human physiology to support language (Lieberman, 1992) (although there are now arguments that the larynx dropped as a consequence of a move to a bipedal stance and that perhaps this was an additional factor *enabling* the evolution of spoken language, a useful exaptation (Aiello, 1996)). At the same time it has been questioned whether the risk of choking is evolutionary significant (Aiello, 2002), other work has pointed to different evolutionary changes that may have been selected for their improvements to the ability to speak clearly (Sanders, 2002).

Additional indirect costs of language are the additional investments that parents must place in children to rear them. In many mammal species, infants have reasonable amounts of autonomy from birth. Not so with human infants who are totally dependent for a considerable period of time after birth, largely due to the need to complete mental development. In order to support the extra energy cost of large brains, a reduction in energy costs elsewhere was necessary and the gut size was reduced necessitating a move to a more nutritious diet (Dunbar, 1996).

Two common fallacies that support the idea of language occurring by accident are the ideas that the brain is a general-purpose mechanism or that brains vary greatly in design between species. These fallacies are refuted by Deacon (1992) who presents evidence that the brain is modular and that no major new structures have occurred to differentiate human and primate brains, despite fairly significant development and larger relative brain size in humans. Changes in brain organisation are similarly limited, but the changes to the brain are not simply a matter of additional development – different areas of the human brain have been relatively more or less well developed than other areas, when compared to other species (Clark et al., 2001).

Lieberman (1992) and Deacon (1992) describe many of the adaptations in brain and body physiology for speech and language.

Evidence from archaeological studies on the evolution of the human brain relating to particular language adaptations is somewhat contentious with alternative explanations or interpretations existing for much of the evidence (Buckley and Steele, 2002). Despite this, the many adaptations and changes required effectively rule out the possibility that human language somehow developed by a single surreptitious accident.

Thus, the argument goes, language could not be the result of some exaptation of a fully formed language organ, as the cost to individuals of some partially formed language ability would be too great to support were it not immediately useful. The number of changes required, including the dietary ones, could not conceivably have occurred unless over some considerable length of time.

Further arguments against the gradual and continuous evolution of language are based on claims that it is not possible to have ‘half-a-grammar’, and that language evolution (perhaps from some proto-language) had to occur in an all-or-nothing discrete, single step (Bickerton, 1984) (an argument now modified to take in two or three such steps (Calvin and Bickerton, 2000)). Burling (2000) contests this and points to different stages of child language to demonstrate that partial grammars can – and do – exist. Despite their limitations compared to full human language, such partial grammars (partial only in being less capable and powerful than modern grammars) are still useful, and would confer an adaptive benefit to their users.

2.4.3 Questioning the Innate Linguistic Ability

While disagreement exists over how innate language ability came about, it is generally agreed that there *is* some innate ability. Which leads us to the next question: what exactly is it that is innate?

Chomsky’s theory states that language is innate in that a number of language principles exist in the minds of children before they begin to learn language – and it is due to these that children are able to learn language and that a UG exists across all human languages.

An extreme innatist view is that the LAD/UG is itself innate, and every rule in the UG is somehow ‘hard-wired’ in the brain. Constraints covering aspects of grammar such as branching order, subadjacency, etc., are all hard-wired and each evolved to

serve its particular linguistic purpose. This appears to be roughly what is proposed by Pinker (1994, Chapter 10).

Some evidence to support the idea that specific rules are biologically determined comes from the discovery of hereditary language disorders (Gopnik and Crago, 1991).

However, this ‘strong’ view of an innate UG is under attack. The evidence of hereditary specific language impairment has been questioned (Bates and Goodman, 1999), as has the degree to which language is determined by mechanisms internal to the brain. For example, (Deacon, 1992) argues that language universals are shaped not only by the structure of the brain and its limitations, but also by pragmatic limitations and peculiarities of the medium used. Such considerations are also mentioned by Dunbar (1996) when he studies size of informal conversation groups, and other factors in human social behaviour. Thus language and its use depends on a host of limitations, including many imposed by the processes of speech production and perception as well as internal physiology and external environment. So not all features of the LAD can be said to have ‘evolved’, and not all are necessarily innate, and coded in the human brain.

Additionally, cultural processes may have significant influence. Kirby (1998) suggests that some language universals may be the result of historical evolution of language in the cultural domain rather than the result of evolution acting on genes. As language is passed from one generation to the next, the ease with which heard utterances can be parsed, and the ease of generating such utterances, may lead to some innately possible grammars not appearing in any human language. In other words, all human languages exist within the space of innately possible languages, but they do not necessarily fill the space. This view is expanded upon in Kirby (1999).

So, evolution in language form may be responsible for some features otherwise assumed to be innate in the Chomskyan LAD. Care must be taken in assuming that any feature of language is due to genetic or phylogenic rather than glossogenic (cultural) evolution.

The innateness of language is also questioned by Elman (1999), who argues that the idea of what it means for something to be innate is itself under-specified. In being ‘innate’, language is somehow constrained by the brain. Representational innateness, where the precise pattern of neural connectivity is pre-specified, is ruled out as the genome is incapable of encoding sufficient information to do this. Elman concludes

that language is innate but emphasises that grammar is not encoded in genomes. Rather, the innate grammar is a consequence of many interactions of genetic expression – in e.g. determining general properties of neurons, local connectivity rules and brain structure – and of ontological development – most importantly in neural development and learning. These last arguments do not rule out strongly-innate language; instead they question *how* language may be innate. This subtle rethinking, from a biologically and genetically innate LAD to an *emergent* LAD is the unifying theme of the papers in the book in which Elman’s paper appears (MacWhinney, 1999).

2.4.4 Why and How? Just So Stories and Grand Theories

Perhaps the most controversial aspect of studies in the evolution of language is the attempt to explain the *reason* for the evolution of language in *Homo sapiens*.

Such attempts are generally contentious and hard, if not impossible, to prove or disprove. A common criticism of work which attempts to explain in detail the process of exactly how language evolved in humans is that the resulting theories are “just so stories” – stories which ultimately cannot be proven because of a lack of evidence, language leaving no direct physical trace in the fossil records. However, by studying the evidence that is there, it is possible to suggest what the evolutionary pressures were that lead to language and, since Chomsky’s exposition of an innate LAD, attempts to explain the origin of language have again become popular.

While not particularly concerned with the question of the functional origins of language, Lieberman (1992) claims the function of language is obvious:

“The contribution to biological fitness is obvious. The close relatives of the hominids who could rapidly communicate *Look out there are two lions behind the rock!* were more likely to survive, as were hominids who could convey the principles of the core and flake toolmaking technique in comprehensible sentences”

(Lieberman, 1992, p23)

However, this argument is flawed. Vervet monkeys can make do with simple warning vocalisations to identify and warn of a variety of predators, why not hominids? To convey the principles of a simple tool making technique, a lesson in the form of a demonstration only would be more effective than one that was given by speech alone. There must be more to the fitness benefits derived from language.

(Armstrong et al., 1994) presents an argument that syntax evolved from signed language in early hominids. One of the main claims for this is that primates are more able to learn to communicate with humans through signing than through vocalisations - that primates have more wilful control over signing than over vocalisation. However, there is a lack of clear evidence that such signing is natural to primates and there is some evidence that it is not (Pinker, 1994, p338-342). Further, apes and chimps do have some physical ability for varied sound production, and are presumably only lacking in mental structures to finely control speech production and, more importantly, the capacity for language be it some form of LAD or whatever (Pinker's 'Language Instinct' perhaps).

The other main claim is that manual signing, by its very nature, includes some elements of syntax - miming an action includes both action and object elements. Such signing is not necessarily at all complex, however, and there is no argument presented claiming that primates naturally use strings or combinations of signs in a way that indicates some elementary grammar.

Other theories, generally not currently favoured, include ones that propose language emerged from singing (for example, Skoyles (2000), and recent evidence linking one of the key regions of the brain for grammar and language with a role in music appreciation (Maess et al., 2001)).

Such theories, positing a particular 'original purpose' or 'original method', are contentious and generally limited in evidence. One alternative theory that is popular is that language evolved out of its usefulness as a social tool (c.f. Dunbar, 1996; Deacon, 1997). This theory is quite broad, and it avoids picking a single, highly specialised 'original function' of language, instead suggesting a range of language functions. As Deacon puts it, "Looking for the adaptive benefits of language is like picking only one dessert in your favorite bakery: there are too many compelling options to choose from" (Deacon, 1997, p377).

Dunbar (1996) suggests that language evolved primarily to allow the formation, and maintain the cohesion, of larger groups than is possible with physical grooming (although clear evidence of early hominid group size is lacking (Buckley and Steele, 2002)). Language can also perform a variety of other useful functions within such populations. Language may be used for grooming and gossip, with many benefits. Some of the benefits of language suggested include:

- Grooming: maintaining relationships between individuals. Physical grooming is generally limited to two participants. Verbal grooming could be used to maintain a relationship between one speaker and around three listeners. This allows individuals to maintain relationships with a greater number of other group members, and ultimately allows the formation of larger groups. Larger groups benefit from mutual support against other groups and predators.
- Gossip: gossip is useful for social cohesion, and transfer of information about others in a social group - which is possibly the most important information that can be known in such a group. Gossip allows an individual to learn about the strength, honesty and reliability of potential rivals and allies without the need for direct observation. This also reduces the time required to form new relationships.
- Second hand information: more generally, a large variety of information about the world can be learned without the need for direct experience and repeated trial and error by every group (Pinker and Bloom, 1992). This knowledge can be complex, extensive and detailed even across generations. Children can learn what is dangerous and what is nutritious without having to taste every mushroom or walk up to every animal. (Also see Cangelosi and Harnad, 2000; Cangelosi et al., 2002)
- Mate selection: language can also be used to help advertise quality, in the search for mates. The need to keep a mate entertained could drive a mental arms-race between males of a species in the quest for females. Those who succeed will have more offspring driving the evolution forward.
- Kin selection: the development of different languages and accents allows easy identification of members of different groups.
- Cheats: members of a group can gain at the expense of others if they can secure the co-operation of others, without contributing themselves. This can be thought of as a form of the Prisoner's Dilemma. Language gives defectors the ability to convince others, and gains co-operators the ability to look for evidence that another individual might not be reliable. This could also lead to an evolutionary arms race, as evolutionary pressure exists to be better able to convince others as well to be able to detect attempts at manipulation (Pinker and Bloom, 1992).

Worden (1998) goes further. Not only did language evolve to serve a variety of social functions but, he argues, it piggy-backed on brain adaptations for dealing with social

situations, an argument similar to (Deacon, 1997). He claims that language ability is based on similar mental structures used for social reasoning. Social situations are structured, complex and open ended, discrete-valued, extended in space and time and dependent on sense data of all modes. This list compares well with the lists provided above for the characteristic features of language. By citing the ‘speed limit’ on evolutionary change, it is argued that it was not possible for speech to emerge in humans given the time since evolutionary divergence from our common ancestor with chimpanzees, unless language was based such on pre-existing structures.

2.5 The Cultural Evolution of Languages

Having surveyed work on the evolution of communication and cooperation and the EoL, one important topic has so far only been mentioned as an aside – that of the eol. People have been long aware that languages change over time, and there is no shortage of quotes from literature about the benefit to be gained by stopping language change, or diatribes on the degeneration of language. For example:

“... nothing would be of greater Use towards the Improvement of Knowledge and Politeness, than some effectual Method for *Correcting, Enlarging, and Ascertaining*, our Language;”

(Swift, 1712)

Or:

“Standard English is the language of English culture at its highest levels as it has developed over the last centuries... This does not mean that speakers of non-standard English cannot be verbally agile within certain areas of discourse, nor that topics traditionally discussed in the standard language are *entirely* barred to them.”

(Marenbon, 1987)

Away from such reactionary views, language change has been the subject of serious study for many years. In this section we provide an overview of some of this work, from a number of different approaches.

2.5.1 Language Change

Change in language is studied at many different levels, the individual, social and historical. Typically studies concentrate on how language change operates at only

one of these levels, and this has led to the emergence of a number of complementary (and potentially contradictory) linguistic disciplines.

Psycholinguistics is concerned with how language works at the individual level. Here language change may be viewed as the result of different functional pressures on language as speakers try to communicate with listeners (e.g. Slobin, 1979).

At the social and group level, studies of language change take place in the fields of socio-linguistics and dialectology. Studies in dialectology show how changes may be spread through populations – or may simply ‘map’ the different language forms that are in use in different population groups. Such studies are generally *synchronic*, attempting to describe the structure of language at a fixed point in time. As a consequence of studying the differences in dialect between different geographical and social groups, it became clear that each individual in a linguistic community uses a unique variant of their dialect. The term *idiolect* was coined to refer to the unique variant of language possessed by any individual in a population.

Sociolinguistics concentrates instead on how individuals actually *use* and change their use of language according to social situation, and on providing explanations for use of particular linguistic variations according to social factors (class, sex, grouping, etc.). See, for example, (Tannen, 1994; Chambers, 1995; Trudgill, 1995) or for a classic study, Labov (1972).

In historical linguistics, much work is done to reconstruct languages no longer spoken, or to trace and explain the changes that may have occurred in the history of a particular language or family of languages (e.g. Trask, 1996). In contrast to sociolinguistics, work in historical linguistics is by its very nature *diachronic*, considering as it does changes in language over periods in time.

While these different approaches are to some degree isolated from one another, they are all interested in language change, how it happens and why. Yet, their different interests and perspectives can lead to very different opinions on what is actually important when describing language differences. For example, in socio-linguistics, which concerns itself with language in society, an individual’s idiosyncratic variations in language use are not considered important (as opposed to systematic variations in language use observed amongst some group of language users):

“Discovering how various “personality factors” interact to make idiolects would probably not repay the effort because they carry almost no social significance”

(Chambers, 1995page 85)

While language change is a phenomenon which occurs at many levels – internal and external to speakers, and over many generations of speakers – the explanations given for language change are almost invariably given at the same level at which it is studied.

2.5.2 Explanations of Language Change

Aware that the type of study affects the conclusions, what explanations of language change exist in the different fields that study it? In attempting to characterise each of a number of distinct fields in only a few sentences condemns me to provide grossly simplistic views – but representative ones, I hope. Sociolinguistics, psycholinguistics and historical linguistics all concern themselves, to a greater or lesser extent, with language change. The diverse methods used and distinct focus of these fields has led to different approaches to explaining language change.

Sociolinguistic studies of language change typically observe differences in language use by different social groups within a geographical area. The spread of change is observed as it moves both *across* and *within* different social groups. Additionally, studies also note how language use is modified according to social situation, and the social benefits of having the appropriate dialect in any social interaction (Chambers, 1995).

A consequence of this is that sociolinguists may view language change as only occurring *because* of these socially functional factors. Why adopt a novel form, if it won't provide some benefit? Regarding sound changes, James Milroy has this to say:

“It must be the case that human beings attach great importance to changes like this: if they did not, then there would be no reason why they should implement them at all”

(Milroy, 1993, page 215)

This view, that language changes and differences have some reason, some functional motivation, is shared by others (Nettle and Dunbar, 1997; Nettle, 1999a). All agree that it is the ability of language differences between individuals to act as a social marker that motivates language change and linguistic diversity. Such a marker would

serve as a badge of group membership, and help prevent outsiders taking advantage of the natural co-operativeness of others (Dunbar, 1996, p169). These arguments will be revisited later.

In contrast, quite different explanations of language change exist in psycholinguistics. Here the emphasis is on functions internal to the speaker or listener, rather than on the social factors surrounding their interactions. For example, Slobin (1979) states four competing functions of language. This competition both fuels and maintains the continued evolution of language. Because of these competing functions, changes gradually occur in languages to optimise one function or another, the competition maintaining language in a dynamic equilibrium. The language functions identified are that language signals should:

- Be clear
- Be processible
- Be quick and easy
- Be expressive

The constant dynamic attempt to maintain equilibrium balances language simplification versus language elaboration. Pressures on language arise because in a conversational exchange the speaker and listener have different goals:

- The speaker wishes language to express meaning clearly, efficiently, effectively and ‘reasonably’ quickly. (Clear, quick and easy, expressive)
- The listener wishes language to allow them to quickly and efficiently retrieve a clear and informative message from speech. (Clear, easy to process)

The competition, and the change it produces, is internal to each speaker and listener, and over some population will cause historical language change.

Historical linguistics, generally being more interested in changes which happen over long periods of time – say, from the English of Chaucer to the English of Shakespeare and on – looks more closely at language itself than at individuals who speak it. What is studied here is generally language as a *system*. How might a change to one part of the system affect the rest of the system?

This view of language as a system gained prominence with the publication of what became known as ‘Grimm’s Law’. In his *Deutsche Grammatik*, Jakob Grimm (1822) explained that correspondences in consonant use across a number of languages were the result of systematic changes from older languages (Latin, Greek and Sanskrit)

(Baugh and Cable, 1978). For example, the ‘*p*’ sound of Latin or Greek could change to the ‘*f*’ sound of modern Germanic languages. A change in the opposite direction – from an ‘*f*’ to a ‘*p*’ would not be possible, however.

Over time this has led to the development of many *Sound laws* which describe ways in which systems of sounds may change over time. Despite the name, these are not laws but more general rules of sound change, for which exceptions often exist (Adamska-Sałaciak, 1997).

These ideas of the sounds of language as a system – and of directional changes in phonology – led to the development of chain-shift theories of sound change. Consider the sounds of language as existing in some auditory space. They may not be evenly distributed through the space. One sound may gradually shift its position in this space, leaving behind a ‘gap’. As the first sound moves away, a second may be ‘pulled’ into this gap, and over the course of time a ‘chain’ of such sound moves forms (King, 1969). Further, explanations that view language as a system are not limited to phonology, but can be extended to other aspects of language (Anderson, 1973).

A limitation of the historical approach is that it often does not explain why languages change at all – rather it provides a framework for studying the histories of changes in languages, and for describing which changes are likely or indeed possible. Viewing language as a system which is itself subject to change loses sight of the mechanisms by which change occurs. Although historical linguistics is well suited to documenting and detailing changes in language use, some of its methods are limited in their ability to explain *why* language changes occur, as the rules and explanations put forward are those that apply to languages as a system – and not to the speakers of the languages (Milroy, 1993).

Recently, however, a number of authors – linguists and others – have made serious attempts to use evolutionary theory to explain language change, and some of these are reviewed next.

2.5.3 Language Change and Evolutionary Theory

The re-application of biological evolutionary theory to explaining (or describing) how languages evolve has a history that can be traced at least as far back as Darwin:

“The formation of different languages and of distinct species, and the proofs that both have been developed through a gradual process, are curiously parallel.”

Darwin, C.R. 1874

However, (neo-)Darwinian theory of evolution is based on many mechanisms that simply do not apply to eol. Languages ‘reproduce’ by transmission from speakers to learners and other speakers. Learners do not apply ‘re-combination’ and ‘mutation’ operators to utterances they hear to form their own language. So how can evolutionary theory be applied to something like language?

2.5.4 Cultural Evolution

This question was asked by Cavalli-Sforza and Feldman (1978), in their attempt to develop a general theory of cultural evolution. Rather than obtaining genetic material from two parents, cultural ‘material’ is received from the community around an individual. The amount of influence held by members of the community is unequally distributed. Community members may have from very strong to negligible influence on another individual. The socially acquired characteristics will be determined by a summation of the influences exerted on the individual. This is expressed mathematically as:

$$X_{i,t+1} = \sum_{j=1}^N w_{ij} X_{j,t} + \varepsilon_i \quad (2-1)$$

This equation is used to determine the traits, X_i , at time $t+1$, of the i th individual in a population of N individuals – based on the trait values of the population at the previous time step. The proportion of $X_{j,t}$ contributed to $X_{i,t+1}$ is determined by w_{ij} , indicative of the amount of influence or contact between the two, with $0 \leq w_{ij} \leq 1$ and $\sum_{j=1}^N w_{ij} = 1$. Finally, ε_i is a random error term.

The equation given is more than strikingly similar to description of a social network used by (Milroy, 1980). A social network is, quite simply, a network describing an individual’s contacts and bonds with others. A portion of a hypothetical network for an individual called ‘Ted’ might look like the network shown in Figure 2.1

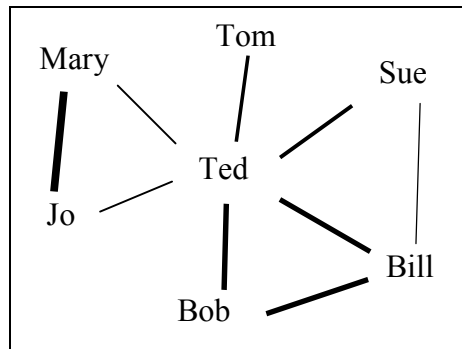


Figure 2.1. A partial social network for an individual called ‘Ted’. Excepting links between members of Ted’s social network, the social networks for the other individuals are not shown.

Each individual is linked to many others in the community around them – through family, friends and work. In some instances an individual might be a member of a close-knit community, with a relatively small number of strong bonds. Other individuals might have more diffuse networks, with many weak bonds. Thus, the social bonds linking individuals vary in their strength. To describe these bonds mathematically we can use Equation 2-1, above: every individual in the community has links with every other individual, varying in strength from zero (no contact) up to one.

While Milroy was not attempting to formulate an *evolutionary* theory of language change, it is interesting to note that social networks map very naturally onto a separately derived theory of cultural evolution.

Cavalli-Sforza and Feldman (1981) present a reworked mathematical model and Niyogi (2002) compares his own computational model of language change with this later model. We will look at these in more detail in Chapter 6 where we study cultural evolution in more depth.

2.5.5 Language Ecology and Species

Two biologically inspired theories of language evolution with much in common are *language as species* and *language ecology*.

In biology, ecology is concerned with the interrelationships between an organism and its environment – which includes other organisms. Language ecology (Haugen, 1971; Mühlhäusler, 1996; Mufwene, 1997) considers human populations as environments within which languages exist. As with its biological counterpart, language ecology considers the initial conditions of an ecology – the starting language and population distribution – and stochastic events to be of importance in the ensuing evolution and

development. Mufwene argues that variations in population sizes and limits on interactions between people (social structure) strongly influence the process of creolisation. Examples demonstrate how these influences can lead to radically divergent eol in situations with what might be considered similar starting conditions.

It is postulated that similar processes are at work in the normal eol: each new generation learns language from scratch based on the existing language around – whether that language is pidgin or an established language.

Within the ecology, Mufwene regards language to be a species; the language of an individual speaker – which will vary at least slightly from all others in the population – is considered to be an organism.

Like animals of the one species, variation is always present, even in stable populations. Sometimes a particular variant may propagate through the population, changing it. At other times a population may remain stable over long periods. Mufwene considers that “the question for historical linguistics is to determine under what ecological conditions... small actions of speakers amount to... change in the communal system” (Mufwene, 1997, p328).

Lass develops a similar view, but – borrowing from the population biology of viruses – promotes the notion of language as a ‘quasi-species’ (Lass, 1997, p375). A quasi-species is a highly variable yet self-stabilising population. Again, the individual within a language quasi-species is a single idiolect. A model of evolution, neutral to what is being evolved, is presented with arguments to show that it applies equally well to language as it does to virus quasi-species. This view makes explicit some aspects often overlooked in objections to language evolution theories that limit the roles of function or intent in eol. I will briefly describe below some of the main points specifically as they apply to language.

The average language of some population may be considered the norm – but this may not be represented by any existing idiolect. The norm is a weighted average of the population, which may be more densely clustered round some point than elsewhere but which may have many outliers on the periphery. Over time the precise shape of the “cloud” of idiolect will change and the position of the norm will shift. Clusters may form away from the norm and these may be favoured and become a new selection-gradient for the language quasi-species. A new norm would be formed inside the, formerly peripheral, cluster. Using traditional linguistic terminology, we could simply say that the language had changed.

Lass coincidentally concludes his argument with the same point as Mufwene:

“What remains then... is to differentiate the ecological conditions under which selection gradients arise”

(Lass, 1997, page 380)

2.5.6 Neutral Evolution

With his work on neutral evolution, Kimura, (1983), showed that it was possible for evolution to occur without any apparent selective forces at work. The possibility that language evolution could be evolutionarily neutral is discussed by Lass as well as by Nettle (1999a), who both refer to Kimura. However, they have very different views on why the neutral evolution of language should, or should not, be sufficient to cause change and diversity in human languages.

Lass’ proposal (Lass, 1997, p. 354) is based on the observation that languages are imperfectly replicating systems, within which elements of linguistic ‘junk’ and other ‘marginal’ features exist. This provides ample room for variation, and allows changes to occur without disrupting the success of communication. That replication is not, and can not be, perfect means that languages will change, regardless of functional benefits.

Nettle (1999a) argues against the neutral evolution of linguistic systems on three points:

1. Random changes would be non-directional and could be expected to cancel each other out, due to an averaging effect.
2. With a neutral model it is difficult to account for diversification without geographical isolation.
3. Structural correlations in many of the world’s languages represents parallel evolution, showing that the path of linguistic diversification is not random.

Thus, Nettle proposes that in order for linguistic evolution to occur without geographical isolation, additional mechanisms are required. As previously mentioned, Nettle argues that the social functions of language are required for the emergence of linguistic diversity (section 2.5.2).

However, Nettle’s first and second points both rely on the equal distribution of individuals, with a uniform likelihood of any one individual interacting with any other. As in the theory of cultural evolution (section 2.5.4), in any group the amount of influence exerted on any one individual by any one of the others will vary

according to a number of factors. This reduces the effect of averaging, and increases the potential for sub-populations to vary from the mean. The different social networks within groups reduces the need for geographical isolation to produce linguistic diversity.

We will revisit the concept of neutral evolution, and arguments for or against it, with regard to language change, in Chapter 7.

2.5.7 Co-evolution of Languages and Man

Coevolution is a term which usually applies to the processes by which two species may evolve, with the evolution of each species influencing or inducing evolutionary change in the other (Stearns and Hoekstra, 2000); it is typically used when referring to the coevolution of predators and prey, or parasites and carriers. A more generalised view of coevolution is that it refers to any two interacting, evolving, systems that influence each other's evolution – say, the different organs of animals of the one species. Like many other evolutionary ideas, coevolution can be invoked to try to improve understanding in areas apart from biology, including the evolution of language. As the ability to use language has evolved, so have the languages that are used (this cultural evolution of languages is discussed further in Chapter 7). As successful use of language became important for survival, this would have in turn put pressure on the evolution of the LAD.

Viewing the LAD (in whatever form it takes) as the biological language organ, the languages used are obviously constrained by the LAD. The LAD shapes the evolution of languages. If it is a coevolutionary system, it is possible that the LAD has in turn evolved in response to the languages that it allows to exist.

2.6 Summary

In this chapter I have provided a short review of a great many topics in linguistic and evolutionary research. In the next chapter I look at computational approaches to evolutionary modelling, and review some of the methodological ground rules that have been developed for conducting such work. These have given rise to ground rules that I have tried to follow in conducting my own work, described in the later chapters.

Chapter 3 Artificial Life: Computational Modelling and a Methodological Approach

Over the course of the last chapter, a broad introduction to language evolution and related problem domains was given. It is the aim of this chapter to provide some detail of the investigative approach used.

3.1 The Artificial Life Approach

The quote from Langton, reproduced in Chapter 1 identifies the essence of much of ALife as the construction of systems in which local interactions of many entities, or *agents*, give rise to an *emergent* and *life-like* behaviour.

As well as forming a definition of ALife, Langton's statement reveals that ALife was even then a diverse field, where the work of different researchers had seemingly little in common. A review of the proceedings from a more recent ALife conference will similarly reveal a diverse range of work.

However, other than the "essence" described by Langton, there is a commonality that can be found in much ALife work. Much of the research tries to improve understanding of phenomena in the real world - whether it is the emergence of life itself, the evolution of co-operation, signalling or language or the dynamics of price-wars in a free market (see, for example, Adami et al., 1998). So despite Langton's assertion that ALife represents "life-as-it-could-be" (Langton, 1991), it can be seen that a major use is as a scientific tool for investigation of the real world.

3.1.1 Agents

The basic element in most ALife models is the agent. In this context, an agent is a single individual in some simulated population. The agents may be very simple and abstract, interacting with one another and/or with a simulated environment. Rules govern the behaviour of agents during interactions and also when, and which, agents interact. The agents themselves may be as simple as simulated 'billiard balls' to more complex organisms that adapt and learn (Holland, 1998).

In theoretical biology, such an approach has been termed *individual-based-modelling* (Grafen, 1991), and a similar approach is also used in work in social science micro-simulation (Gilbert and Troitzsch, 1999).

The simulated populations may consist of handfuls of agents, or many thousands of agents, interacting according to simple rules. This is in contrast to traditional mathematical approaches where entire populations may be modelled by means of a system of complex mathematical equations (such as in traditional dynamical systems, see Hofbauer and Sigmund (1988)).

It should be noted that, for some, the term agent has a different, and more strictly defined, meaning than is used here. Some definitions require that an agent is embodied in an environment with which it is able to interact and/or that an agent must possess its own (however modelled) beliefs, desires and intents (Wooldridge and Jennings, 1995).

3.1.2 Emergence

Where a phenomenon at one level arises as the result of processes that occur at another lower level, the phenomenon can be described as an emergent one. In physics, a classic example is temperature. Individual atoms and molecules do not possess temperature – it emerges from the interactions of many different atoms and molecules.

In ALife, the classic example is ‘Boids’ (Reynolds, 1987), a model of flocking birds. In this model, individual boids follow simple rules to avoid collisions, to try to match velocity with other boids and to try to stay near the centre of the flock. From the interactions of many boids following these rules an emergent, and life-like, flocking behaviour emerges.

The emergent phenomena itself may be subject to “macro-laws” – laws which describe the emergent behaviour at a higher level than that of the interactions which give rise to it (Holland, 1998, p 224) – temperature again providing a good example.

3.1.3 Life-Like

The purpose of ALife modelling is usually either to allow future events to be predicted or to provide better understanding of existing phenomena. In this thesis I will generally adopt the latter, as ALife models – due to their nature, where initial conditions and the stochastic events can radically affect quantitative results – are not always well suited to making predictions other than in the most general terms. Whatever the purpose, the method of obtaining results is the same. Having initialised a population of agents and their environment, set parameter values, and iterated the

simulation for some period, the results of the simulation must then be compared to observations of the real world.

While it is possible that numerical and statistical analysis of real and simulated data will allow a direct comparison, it may be the case that some amount of interpretation is required. In such a case, apparent or qualitative similarities between model and world are identified – as in the flocking patterns of Reynold’s boids.

With comparable results between the simulation and the real world, it can be postulated that the processes which produced the phenomenon of interest in the model are essentially the same as those at work to produce the comparable real-world phenomenon. In many cases it is hoped that this provides some degree of explanation for the phenomenon. Note, however, that a similar end result does not necessarily mean that the same processes are at work, as we shall later see.

3.2 Methodologies for Applying Artificial Life

Reviewing ALife research which aims to model the real world, it is apparent that there has been a lack of discipline and rigor, and this has been noted by a number of authors – some of whose suggestions for improvement are reviewed in this chapter. Indeed, to some prominent researchers in the field of Evolutionary Computation, which focuses on the scientific applications and understanding of evolutionary algorithms, the related field of ALife is seen as an area characterised by poor quality research, where researchers simply play around with models and in which an “anything goes” attitude is prevalent (Muehlenbein 1998, personal comments).

This compares poorly to established sciences, where higher standards of rigour and use of appropriate scientific methods are expected. In the social, physical, and life sciences criticism of the methods by which a finding has been reached will substantially undermine confidence in the findings. The requirement that suitable methods are applied to research before the research can be accepted generates interest in, and awareness of, methodology and many volumes have been published on suitable methodologies for research and experimentation. For example, a search of the University of Paisley library (a rather small university library) catalogue for texts with the words “method” or “methodology” in the title lists over 700 separate titles. In contrast, a small, but growing, number of papers have been published providing guiding principles and heuristics to apply when conducting ALife research.

Before using ALife methods to investigate the evolution of language in the following chapters, in this chapter I will review current ideas on ALife methodology and good practice. In the remainder of this thesis I will endeavour to satisfy these standards.

3.2.1 Science without method?

In the first issue of the *Artificial Life* journal, the paper “Artificial Life as a Tool for Biological Inquiry” (Taylor and Jefferson, 1994) was published. The potential of ALife for scientific inquiry in the biological sciences is discussed and some problems in biology where ALife may be usefully applied are presented. However, there is no mention of methodology or recommendations on how to conduct such research.

The proceedings of the early Artificial Life workshops produced a lot of important work establishing the field, but very little groundwork on methodological issues. Instead the emphasis was on showing that ALife could be applied to a varied range of problems, helping to find answers to complex problems. The advantages of ALife models over traditional mathematical modelling techniques was emphasised at the expense of cautionary notes on the limitations (e.g. Taylor and Jefferson, 1994).

Part of the problem is that ALife is often seen as being a method in its own right, or as a completely new science to which established scientific practice need not apply. Noble (1997) questions such views, asking if ALife is a science, and, if it is, then what does it study and what standards should be applied to its practice. His conclusion is that ALife is *more* a method than a science, i.e. a scientific approach that can be applied to a target discipline rather than a standalone discipline itself (although, ultimately such distinctions are problematic – consider the case of statistics as a parallel). As we shall see, a key problem is that, as a new *method* for conducting scientific inquiry, ALife has been lacking in developed *methodologies* to guide its application.

Where traditional mathematical models need to make more general assumptions about the activity of individuals within populations to make the maths tractable, ALife models can leverage modern computational power to allow models to be built which make fewer assumptions and simplifications. Yet, other problems are introduced by the use of computational models.

An example of the problems that can occur is presented in Noble and Cliff (1996), itself a critique of an earlier work on the evolution of communication, (MacLennan

and Burghardt, 1994). In the critique, Noble and Cliff are compelled to present some warnings regarding the use of ALife simulation. These are:

- to beware of counter-intuitive results due to the conditions used in the experiment
- that simulations with unnecessary complexity may show nothing
- and that if an arbitrary decision about the simulation implementation or design influences the results, then there will be problems relating the results to the real world.

Di Paolo (1999b, chapter 4) identifies the core problem with the use of computational models in research as being ensuring that the conclusions drawn from experimentation with a synthetic model are realistic, apply to the real world and truly answer the questions asked of them.

These types of problems are noted by Gilbert and Troitzsch (1999) who discuss the use of simulation as a method in social science. The problem of determining whether results could be due to initial conditions or even some programming error in building the model is one of verification; that of ensuring that comparisons with the real-world hold true, validation.

In the next section we review one of the first papers to address this problem, before proceeding to look in more detail at the ALife research life-cycle and the impact of the verification and validation problems on experimental design, implementation and explanation.

3.3 Improved Methodologies and Practices

3.3.1 A Methodology Emerges

Perhaps the first attempt to review and criticise the scientific methods employed, or not employed, in ALife research is presented by Miller (1995). Miller takes a critical look at the use of ALife as a tool for theoretical biology, noting many pitfalls and hazards for the unwary.

Miller first reviews some failings of computer science in general when it attempts ‘real science’ (sciences based on pre-existing natural phenomena). Miller claims that whenever computer scientists develop a new field claiming to be real science the same problems recur: poor scholarship, with interdisciplinary blindness; poor research methods, with a lack of rigorous hypothesis testing, controlled experimentation and statistical analysis; poor analysis of results, refusing to recognise

failures and wildly over-generalising successes; and finally, poor follow through, leaping from research fad to fad, failing to replicate results and extend findings towards a conceptually integrated discipline.

As redress, Miller suggests six methodological heuristics for successful ALife, and these are listed below.

Miller's Six Heuristics

- 1) Identify a known and unsolved problem that can be addressed using simulation. Miller notes that this may be more difficult than it sounds. Biology is a mature and successful science, and most problems a casual reader may encounter have large bodies of work behind them, or just about to be presented.
- 2) Collaborate with real biologists who have already worked on the problem. ALife work may not present much gain to biologists, and may present some risk. Sufficient biological knowledge is required to earn their support and confidence.
- 3) Do a thorough scholarly review of the current biology literature relevant to the problem. This can avoid effort being expended on solved problems.
- 4) Develop a well-targeted simulation that extends current biological models and yields directly comparable results. For example, mathematical models in biology may make strong and unrealistic assumptions to make the maths tractable. Such models are weak at coping with phenomena like complex phenotypes, flexible behaviour, and co-evolution. Taking such models and relaxing the assumptions one at a time is a powerful technique and the results can then be directly compared to those of the formal model.
- 5) Explore cause and effect in the simulation by running comparisons across different conditions. Rather than constantly moving on to new and better simulations, perform strong and thorough analysis of the results from simulations under numerous conditions. Vary independent variables and observe effects on dependent variables.
- 6) Publish the results in biology journals, subject to peer review by real biologists. If the work is unacceptable, then it is likely flawed. Miller views journals within ALife as playgrounds in which ideas and methods can grow before venturing outside. Miller suggests an ideal ALife project lifecycle in which work is iterated over several years with publications moving from conferences to the dedicated journals and finally to mid and then high level journals within biology.

Miller concludes that ALife research will only be as good as ALife research methods, though the prospects are exciting. The comments presented in this paper are a sobering reminder that ALife work purporting to illuminate some aspect of evolution must be presented within the context of the other work that already exists. In particular, Miller warns ALife researchers not to believe that they can enter a well-established field of science and solve all of the open problems with a few simple simulations - or to act as if they believe it.

Overall, these heuristics are helpful, focussing on the particular failure of much of ALife research to properly apply itself to its supposed target problem and domain. The advice here is to select a particular problem, and to ensure that the model does indeed address the problem. While acceptance by peers from the target domain is a worthwhile objective, there can be resistance to accepting for publication work based on computer modelling rather than the more traditional methods used in a particular academic field (for a particular example, see Axelrod, 1997a, chapter 7), and in some cases this may make Miller's final heuristic a difficult one to successfully follow.

3.3.2 Not Just Biology

While many of Miller's points are well observed, the heuristics for guiding research have themselves been reviewed with some criticisms. Some of these are reviewed here, together with other works which provide suggestions for good practice in ALife research.

One obvious point is to ask why should ALife be limited to just research in Theoretical Biology? That ALife has much to offer researchers in other fields of science is noted by both Noble (1997) and Di Paolo (1996), this view easily supported by the large body of work using ALife techniques in subject areas other than biology.

Yet many of Miller's statements could apply equally to those using ALife for linguistics or social sciences, etc., simply by substituting the name of the appropriate science where "biology" appears in Miller's heuristics. This should allow the heuristics to be generalised to apply to a broader range of ALife work.

3.3.3 ALife for Novel Models and Theories

Miller's fourth heuristic comes in for heavier criticism, however. In an unrelated paper Hurford (1996) states that the act of building a computational model can help

find gaps or hidden assumptions in theories, and the running of such models can help improve the detail and check for internal consistency. This view is similar to Miller's point 4, but only requires that a theory exist, not that a formal mathematical model already exists to support the theory. This is an important observation as in social sciences and humanities theories are much less likely to have formal mathematical expressions than they are in the physical and mathematical sciences.

Di Paolo (1996) finds the heuristics more limiting than necessary, and criticises the fourth in particular. He accepts that work following this rule may enrich current models, providing new answers that are too hard, or even impossible, to obtain using traditional analytical methods. However, he questions whether this heuristic may present its own range of problems.

The principal problem identified is that of inheriting implicit assumptions. An ALife model that extends an existing model will have many of the methodological and philosophical assumptions of the existing model. These may be implicit points simply accepted without question in the original domain. Given that in biology there exist a number of unresolved controversies and debates, an ALife model developed from an existing biological model may depend on factors that are not generally agreed upon. In contrast, work that is not developed directly from an existing model may have the potential to help resolve open debates. As we noted above, however, developing a new model can cause problems with the model's ultimate acceptance.

Computer explorations where no prior formal model exists requires first building such a model, a step avoided when using Miller's guidelines. Di Paolo points to physicists who believe that some phenomena can be better modelled computationally than they can mathematically, supporting the principle that direct modelling can be good practice.

Further, verbal arguments are often used where systems involve complex processes that are hard to formalise mathematically. Computer models can be built to represent verbal models, and simulations used to help defend or attack such models. This makes the point that existing non-mathematical models may also be used as the basis for ALife experimentation.

Further, restricting ALife researchers from developing new models and theories simply hands this work over to other researchers, without a good reason for preventing ALife researchers from participating in this work (Noble, 1997). There is, however, no point in advancing a novel ALife-inspired theory to explain some

phenomenon when a better theory already exists – a reminder of the importance of properly researching the problem domain (Miller’s point 3).

Di Paolo (1999b, chapter 4) promotes the development of a number of different formal models (such as mathematical and computational) of some phenomenon, and then comparing the results of the different models to improve understanding. This can help highlight which conclusions depend on particular models and which assumptions and features cause differences in the results. This helps to question and assess the assumptions themselves, providing greater insight into the problem than otherwise possible.

These arguments all hold that ALife presents a valid approach for both (re-)evaluating existing theories and for developing new theories, with the computational models themselves providing demonstrations of the new ideas. There is no compelling argument binding the ALife researcher to working with models derived from existing formal models – although such models are useful for comparison.

3.4 ALife in Practice

Accepting the development of ALife models as an alternative to statistical or mathematical modelling, are there significant differences in how research should be conducted? Or given the apparent gulf between the approaches, are there any significant similarities?

3.4.1 ALife Research vs ‘Real’ Research

Gilbert and Troitzsch (1999) shows that the general approach when using simulation as a research method is the same as when using traditional modelling methods. Traditional modelling proceeds in a number of stages. Observations are made and some data collected. A theory is proposed and a mathematical model developed. For a number of initial states and parameter settings the model is used to derive future states. Comparisons between predicted and collected data can be used to support the validity of the model. The same stages exist when developing a computational model. Instead of comparing predictions derived from a statistical model to collected data, some similarity is searched for between the data and the results of simulation runs. Differences will exist in the details of each stage, but the overall process is essentially unchanged.

Di Paolo (1999b, chapter 4) also considers how research should be conducted using computer models. He describes a possible way to integrate simulations into a scientific project, to help assure scientific integrity without limiting the potential of the simulations. Three distinct phases of research are identified, two for working with a simulation, after the initial model has been constructed, and the third for relating the results to the real world:

- 1) An Exploratory Phase. During this stage, different cases of interest are run, observables defined and patterns explored. Reformulation is performed as required (by unexpected results).
- 2) Experimental Phase. Hypotheses are generated, and further simulations run to test the hypotheses.
- 3) Explanatory Phase. The observations must be related to the theories and hypotheses about the natural phenomena in question.

Di Paolo also notes that this methodology has been used before, and is not unique to ALife. He also makes clear that this is only one possible way of attempting to use ALife models according to scientific principles.

As theories can follow on after the model has been constructed, it is also important to consider guidelines on how a model should be constructed.

3.4.2 Principles of ALife Model Building

A couple of principles for ALife model building are agreed upon by almost all of the authors cited in this chapter, and others beyond. These are the need for *minimal models* and the need to make any assumptions captured in a model explicit.

“Models should be as simple as possible but no simpler!”

(Doran, 1996, p382)

A minimal model is one which models all of the mechanisms required by the theory and no more (or as little extra as possible). There are a number of reasons for preferring minimal models.

First, given two explanations for some phenomenon the simpler explanation, requiring less assumptions, is often the more likely (the principle commonly known as Occam’s Razor). Second, attempts to capture more detail than strictly necessary, perhaps to increase ‘realism’ in the model, may detract from the simulation. It may make it harder to observe what is happening in the simulation or to determine which of the increased number of factors modelled are the key factors in giving rise to a

certain result. If that is not a problem, there might still be an increased difficulty in verifying results. Additionally, one undesirable result of using a non-minimal model might be an increased computational load – resulting in simulations taking longer to run.

Noble (1997) describes a minimal model as being one that captures all *and only* the intended assumptions. This highlights another point, that there is a need to ensure that any assumptions built into the model are made explicit. Such assumptions have to be included in the theory, or otherwise accounted for in any explanation of the model. The subsequent discovery of hidden, unacknowledged or unknown, assumptions requires additional effort in verifying and validating models (see 3.5 and 3.6).

3.4.3 Adapting Methods

Taylor (1998) discusses specific issues that relate to using ALife to model the evolution and origin of life itself. It is observed that in this specific field of ALife endeavour many existing models are lacking in explicit statements of what exactly is being modelled, that some definition of life is required, and that explicit lists of assumptions are often also missing. Taylor has identified particular problems in using ALife to tackle a particular problem, and has suggested some particular solutions. This itself is an application of a more general methodological rule, which could be stated as:

- Appropriate methodology for using ALife to investigate some real-world phenomena will, to some degree, depend on the phenomena under investigation and the scientific discipline which studies it.

This conflicts with all of the attempts described above at deriving general principles for conducting ALife research, but does not strictly contradict them. Rather, it points out that the general guidelines are insufficient for any individual research program. An important step will always be to adapt the general methodologies for ALife based research to the particular problem domain as a single prescriptive methodology will not be applicable to all areas to which an ALife approach may be applied.

While it would be ideal for any researcher to have a fully developed methodology at the beginning of any program of research, this is perhaps an unrealistic expectation. More likely, particularly in the case of doctoral research, a methodology may co-evolve with the research itself.

3.4.4 Implementing an ALife model

The discussion on methodology presented so far has ignored the details surrounding the actual implementation of an ALife model, which is now briefly considered. Both ALife and traditional formal models require some amount of abstraction. To build a mathematical model requires the researcher to abstract the interactions and processes that occur in a given population into a single set of equations. Values may be ascribed to different variables and the result is expressed again in mathematical terms.

In sharp contrast, the ALife modeller has to write a computer program, which contains within it a simulated population. This is a particular weakness of ALife – it depends ultimately on writing a computer program, and programming is often considered less of a science and more of an art (leading Donald Knuth to title his seminal work ‘The Art of Computer Programming’ (Knuth, 1969)). Many volumes are dedicated to methods that should be employed in order to produce computer programs that do what they are intended to do, and do so without error. Attempts have also been made to formalize the creation of computer programs, but with little success outside of safety-critical systems (Storey, 1996).

And yet there are almost no guidelines – a few are to be found in an appendix of (Epstein and Axtell, 1996) – on this aspect of ALife research, perhaps because the programming exercise is not seen as part of the research itself, rather than as a tiresome task to be done before the research begins.

The first task of the modeller should be to develop a description of the model that they will implement – the agents or actors that comprise the population and the rules which govern their interaction. In ‘normal’ software development, such a description and the methods by which it is developed are of paramount importance, with detailed algorithms and/or system diagrams being required before programming begins.

The software design and implementation paradigms, most suited to ALife modelling, are Object-Oriented Design and Programming, hereafter OOD and OOP. In an OOD, the design is based around identifying the objects that exist in the system, their properties, known as *attributes*, and the activities which they can perform, or *methods* (e.g. Bennet et al., 2001). With OOD and OOP the process of designing and building an ALife model is based around providing a clear description of what the agents and any other objects in the model are, what their properties are, and what actions they can take to interact with other objects or agents. While many ALife models, including

the ones described in this thesis, are developed using OOP, further work is required on the possible benefits OOD can bring to those wishing to implement ALife models.

3.5 Experimentation

Once a model has been built, the visible part of ALife research begins – running simulations under different conditions and observing what happens within the model as a result. Unsurprisingly there are a number of problems particular to ALife experimentation and a number of useful methodological rules to help. Here we concentrate on the considerations of whether a model and its results are valid and on some of the problems facing attempts to verify the results of ALife work. We briefly review some arguments over the place of, and responses to, surprise results in ALife.

3.5.1 Validation

For any abstract model of reality, an important question is whether it is a valid model. Are the principles on which the model is founded sound, and are the abstractions that have been made reasonable? Validation of a model is the key step of certifying that the model and results together provide a legitimate demonstration of corresponding processes and outcomes that occur in the real world.

The principles and assumptions the model is based on must be correct for the model to be valid. So, for example, consider Reynolds' boids (Reynolds, 1987). While the flocking behaviours produced are very life-like, the rules themselves are less so – requiring that each boid is able to compute the position of the centre of the flock at any time (Noble, 1997). It is extremely unlikely that birds are capable of performing such calculations. Without ruling out the possibility that birds do something *similar*, perhaps based on a gross estimation, it is clear that the rules that govern the behaviour of individual boids and those that govern birds flying in flocks are not the same. The results of the model are convincing, but the model is based on an invalid representation of how birds act.

Ensuring that a model is valid is, like many of the guidelines in this chapter not a problem only for ALife research (for example, the validity of the assumption of rational behaviour underpinning a great deal of research in economics has been questioned (Ormerod, 1994)). What makes ensuring that a model is valid more of a concern in ALife is that each model may be introducing its own sets of assumptions and abstractions, as each model may have its own unique agent design and

implementation. These will then be less subject to thorough peer review than shared assumptions developed over time and used collectively by different researchers working in, say, a statistical or traditional game theoretic framework.

However ensuring that an ALife model is based on sound principles does not ensure a correctly working model – that the model actually implements the theory, and that individual agents do what the experimenter believes they do has to be verified.

3.5.2 Verification

There are a number of sources of error that can affect any simulation work, and the results gained from working with a simulation. At best, they act to reduce confidence that the results gained are the correct ones and are not due to the errors present. At worst these might significantly alter the results gained from the simulation.

One obvious source of possible error in a simulation is the presence of a programming error, or bug. With iterative simulations it is difficult to ensure the eradication of bugs, which may not affect the results in an obvious way, such as causing a program to crash, but may rather lead to results which would otherwise not arise. Heralded results may not be a consequence of the intended model, but of an error in its implementation.

The intermediate design stage, linking the expounded theory and abstract model to the implemented model can also introduce errors – or where the design stage is absent, in the direct coding of the abstract model. During this work practical details of the implementation have to be finalised and implemented. Some decisions here will be, in effect, additional, hidden, assumptions. Such arbitrary implementation features of a simulation model can easily affect the results (Noble and Cliff, 1996). Results observed again may not be a consequence of the supposed theory underlying the model, but of unintended effects of some feature of the implementation.

One way to address this problem is through the replication of experiments (Noble, 1997). An implementation by some third-party of a model based only on a published description is unlikely to replicate the same bugs, arbitrary design decisions or hidden assumptions. A replicated result then increases confidence in the published model.

Unfortunately, within a single programme of research, such as a thesis like this, replication cannot be relied upon as one author, or group, is performing the replication. An alternative is to compare the model against other related models, in a search for results that are supportive. This may be possible where different aspects of

the same phenomenon are studied by different researchers. Models will vary, but will have a number of common features. These may be enough to allow comparisons to be made between different models. An example of this will be seen later in this thesis, where two quite different models are used, and their results compared qualitatively (Chapters 6 and 7).

A simpler problem to deal with is the one caused by simulation artefacts that can result from particular combinations of parameter values. Selected values for different parameters can interact in an unexpected manner, perhaps leading to unusual or exceptional results. Some parameters used in a model may be derived from existing observations, and some from theory and many may be arbitrary experimental parameters used in the running of the model. The results of the experiment should reflect the theory not the arbitrary design decisions, and the influence of parameter settings should be minimised. With simulation models it is easily possible to repeat an experiment with different parameter settings, however with a potentially infinite range of parameter values, exhaustive testing is not possible. The solution here is not to test that the same qualitative result occurs for *all* possible parameter values, but that the result gained is *robust* – i.e. the result is not a chance product of a particular combination of parameter values, but is seen over a reasonable range of parameter values.

ALife models generally aim to show how stochastic processes can lead to predictable results, and implementations of this will rely on having the computer provide the simulation with streams of random numbers. Unfortunately, the nature of random number generation in computer programs can be considered problematic. In short, the random number generators are only *pseudo-random*, aiming to give the appearance of randomness (Gilbert and Troitzsch, 1999, Appendix C). The sequences are in fact pre-determined. Using a particular ‘seed’ value, the exact same sequence of ‘random’ numbers can be generated over and over again. To compensate for this each simulation should be run several times with different seed values for every parameter combination used.

Verification problems do exist in other scientific disciplines wherever measurements are made (e.g. Frankfort-Nachmias and Nachmias, 1992), but the additional problems of working with simulations are quite significant. The key solution to this is that wherever possible researchers should seek some form of additional confirmation for existing results. Additional models created by the same researcher may also carry the

same hidden assumptions that exist in the researcher's previous models, and so this should ideally come from the work of other researchers, possibly from attempts to reproduce the same model.

3.5.3 The (Un-)Importance of Surprise

Note has to be made at this point on the appropriate role and response to surprise results in ALife. Some authors place particular value on surprise results in emergent systems (for an example, see Epstein and Axtell, 1996). Some go so far as to include a requirement for a result to be surprising in their definition of *emergence* – and state that where a result is no longer surprising (after a model has been studied, and understanding gained of how particular results come about) the phenomenon is no longer truly emergent (Ronald et al., 1999). Such a 'moving' definition of what is, and what is not, emergent is controversial and unsatisfactory. Indeed, it would rule that classical examples of emergence, such as temperature (Section 3.1.2), be no longer considered emergent at all.

Bullock (1997) argues that without a prior hypothesis of the system and emergent results, surprise results are of little interest – some prior idea of what the results should be is required. Getting the results we expect will, far from being useless, confirm our theory. Obtaining unexpected results, where some parallel between the artificial and the real results can be noted, remains useful and requires further investigation to turn an incidental observation into a scientific one (Di Paolo, 1999b). This might require reformulation of the underlying theory and additional work refining or redeveloping models.

Thus for researchers, progressing through Di Paolo's three phases of research, a surprise result during the experimentation phase may signal a return to exploratory work. While the research can be split into three phases the research effort may not necessarily proceed automatically from one stage to the next, Figure 3.1.

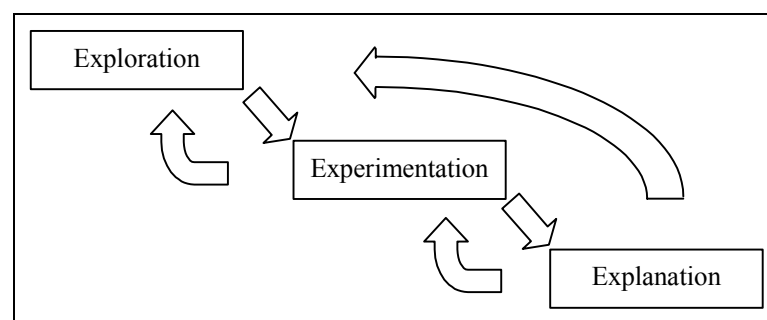


Figure 3.1 In practice progress through the three stages of research identified by Di Paolo may not be straight forward.

3.6 Explanation

In ALife research the importance of explaining results is sometimes overshadowed by the experimentation. Models can be described, detailed and dissected with little reference to the external phenomenon that they are supposed to represent. That it is necessary to detail the model and results in sufficient depth to allow replication can increase pressure on researchers to concentrate on this aspect of research to the detriment of the explanation which is as important.

In explaining the results, the goal is to provide convincing arguments that the model, the results and, most importantly, the theory are all relevant to the target phenomenon. The explanation is required to tie the artificial model and the results obtained with it, whether precise numerical values or general patterns, back to the real world. If the results do not feed back to the real world then the experiments have not been ‘science’, and can have no bearing on evaluating the proposed theory (Noble, 1997), and with poorly explained results an ALife model is meaningless even where care has been taken to verify the results and to replicate the model.

This is somewhat in opposition to Langton’s frequently quoted assertion that ALife represents “Life-as-it-could-be” – we are limiting ourselves to models that help us explain life-as-it-is instead. While this has been a topic of some debate, it is generally the case that ALife research is more devoted to understanding, or at least replicating, real world phenomena than it is to more speculative work. The future of ALife might see greater differentiation, and possible redefinition, but acknowledgement must yet be made of current efforts in using ALife methods in more creative and artistic endeavours (Dautenhahn and Nehaniv, 1999). However, for the purpose of this thesis, we side with those who attempt to use ALife as a scientific tool for understanding the real world. Results of simulations will therefore require comparison with other scientific observations.

(Additionally, note that the field of Evolutionary Computation (Mitchell, 1999) is one in which similar methods to ALife are used but without any attempt to tie models to natural systems. Instead, it is the ability of artificial evolution to act as a means of optimisation or the general underlying theories of evolutionary systems that are of interest.)

For those working in the ‘Scientific ALife’ mould, as described above, there are a number of recommendations to be made about how to explain results, and how to evaluate the success of the model and theory.

3.6.1 Model and Theory

In ALife then, the results of simulation should not be viewed as significant research results in and of themselves, but as a means to evaluating the theory which the model has been built to demonstrate.

Statistical analysis of results may reveal many aspects of the model as a system. If the model has been built to demonstrate a theory, such analysis is likely to be useful, but should not take precedence over explaining the relevance of the model. As noted by Noble (1997) the use of appropriate experimental method and of statistics to provide analysis of results are recommended, but using these does not automatically make an ALife experiment ‘science’. The use of statistics to describe the model should be directed by the goals analysis, which may be one or more of the following:

- To describe the model output in sufficient detail for replication.
- To provide a rigorous description of the system under all possible conditions, including key parameter values causing qualitative changes in system output.
- To detail the output as required for a follow on argument that the results are valid and correlated with those observed in the target and so support the proposed theory.

In most ALife work, the first and third goals will be important, the second goal being important for more theoretical work on complex systems. So, being confident that the simulation model is an accurate implementation of the theoretical model, the results obtained are now the measure by which the theory can be tested.

3.6.2 Validation of Results

Armed with a plausible theory and model, related to the natural process under investigation, it is also necessary to show that results obtained from an ALife model are valid, based on some comparison of the results from the model against what is observed in the real world.

In many cases it will be patterns that emerge in the model which are compared against patterns that occur in real life. The emergence of clusters of agents under differing conditions is a good example of this (e.g. Epstein and Axtell, 1996). Such comparisons are qualitative measures – similar *types* of results are observed in both

systems. While unable to provide precise numerical answers, qualitatively similar results may provide evidence that the types of processes that occur in the model are not inaccurate representations of the real world processes.

Comparing patterns is a subjective measure – the researcher is a subjective observer, interpreting the results and deciding when similarity is close enough to support the theory. Due to this, there is an important need for explanatory argument in detailing the outcome of research. Indeed, any explanation of results will rely on some argument that the results of the computational model actually relate to the real world, and that the abstracted agents in the model interact and behave in a manner that represents the phenomenon being investigated.

In some cases it may be that numerical or statistical measures may allow a direct comparison of the model to data taken from the real world. Quantitative results will often be obtained where the model is intended not so much for explanation but for prediction, in the tradition of dynamic models like Forrester's world population dynamics model (Forrester, 1971). From a set of quantitative results it should be possible to check predictions against real data as part of validation. But are quantitative results possible, and how good are they as indicators of the soundness of the underlying theory? This is one of the aspects to be considered when thinking of the potential limitations of ALife.

3.7 Limitations Of Artificial Life

This chapter has largely concerned itself with how ALife might be applied to scientific investigation, and the appropriate methods for this. Some limitations of the ALife method also have to be recognised.

3.7.1 ALife vs. Mathematical Models

Di Paolo (1999b, Ch. 4) points out that mathematical models have some strengths that are not shared by computational models. Primarily, mathematics is both a means to explore a model *and* a means of communicating the details of the model. Mathematics is a language shared the world over, and a mathematical model can be read by anyone with the appropriate mathematical skills.

Conversely, the inner workings of computational models are generally hidden, the code unpublished. This deficit is clearly one that has to be considered when building software models, and effort must be spent to make the implementation as transparent

as possible. Detailed algorithms should be published and Di Paolo argues that the code that implements a model should be publicly available, presumably at least on request. Even with publicly available code, the separation between the description of a model and its implementation is unfortunate, but unavoidable. Further, the complexity of many programming languages can make source code nearly impenetrable to understand. The importance of verifying models has been previously mentioned, but without a guarantee of accurate translation from theoretical to simulated model this problem will always remain.

Another weakness of a simulation is that it is not necessarily clear what is happening in any given simulation to produce different results (Di Paolo, 1999b, p.68). By means of systematic relaxation of assumptions (Miller's fifth heuristic) it is possible to make strong guesses, but some uncertainty may remain.

3.7.2 ALife Models as Proof

Now consider some limitations of ALife regarding its ability to prove theories. Were this thesis, or its author, of a more philosophical bent then the issues here could certainly have been given a much more thorough treatment – instead the goal here is simply to illustrate some shortcomings of the ALife approach.

Karl Popper's principle of falsifiability (summarised in Thornton, 1997) holds that a theory can never be proven – but that theories may be disproven. Without an infinity of observations how can it be known that B is true for all cases of A? According to a Popperian view, the fact that it has not yet been possible to disprove something is as close to proof as it is possible to get, and this principle is applied in much of contemporary scientific practise. How does ALife fit into this, and can ALife models be used as a form of 'proof'? Obviously it would need to be shown that the model is valid, as already discussed, before the model could even be considered for its role in disproving a theory.

Muller (1996) uses the principle of falsifiability in his arguments for proof by parameter optimisation. If it can be shown that even for the worst possible parameters in some simulation that the predicted result is still obtained then this is some way towards showing that the theory cannot be disproven. This may entail performing the simulation a great number of times for many different parameter combinations – 'good' ones as well as 'very bad'.

Practically, this method may not work. For stochastic simulations, any set of results are due to random processes and hence only one set of observations out of an infinity, so introduces the question of how many simulations have to be run before an attempt at falsifying a theory is considered satisfactory. This could tie up infeasible amounts of time with possibly little gain. It has already been observed in the notes on obtaining quantitative results (Section 3.6.3) that ranges of parameters may give significantly different results. Hence, the worst possible parameters may well lead to results that do not support the theory at all.

In some cases we may be justified in expecting *reasonable* values of parameters, as extreme parameters can actually represent quite different conditions and rules. For example, a parameter might represent the likelihood of some interaction occurring. Extreme values of such a parameter would clearly represent quite different conditions, and would lead to very different outcomes. Rather than test for the worst possible parameter values we should, in many cases, test across ranges of parameters to ensure that a particular result is robust. Attention should still be paid to results for extreme cases, but it may be obvious that the extreme case is one that is not possible for the real world agents and system.

Rather than ‘proof’, what is seen most often in ALife work are *demonstrations* that under certain conditions a certain result is observed. In some cases this may be useful in disproving theories: where a theory claims that some rule is required for a particular result, a model that obtains the same result but without the rule demonstrates that the rule is not required. Even in this case it does not *prove* that the agents in the real world don’t follow the rule – just that it is not required. However, it still may be that an ALife model is better as a test-bed for falsification rather than proof. Again, the experiments detailed in Chapter 6 attempt to disprove existing theories.

3.7.3 ALife Models as Evidence

Alternatively, the results of ALife models may be considered not as proof but as contributory evidence in the construction of a scientific argument (Toulmin, 1958). Rather than present the results of a simulation with a flourish and claim that it conclusively proves a theory, the results may be used in conjunction with evidence and observations of the real world phenomena in building a convincing argument to support a theory.

The known difficulties in validating and verifying ALife work are less problematic when the ALife models are no longer the sole support for some theory – a particular model is just one of many pieces of evidence. To use this approach successfully further emphasises Miller's demands for thoroughly researching the target domain and/or for working with researchers in that domain.

This approach is particularly suitable where there are existing theories that are being explored with the ALife models. Arguments will have already been formed, with supporting evidence, for the different theories. The outcome of the ALife research will ideally be further evidence to support one theory as well as evidence to weaken the support for other theories (this also, obviously, relates to the use of ALife in disproving theories).

3.7.4 ALife and Quantitative Results

As noted above, in many cases the results of simulation models are compared qualitatively to real world observations. The researcher is asking if, for a particular set of rules and parameters, the same *sort* of results are observed in the simulation as in the real world. A different view is put forward by Muller (1996), who argues that quantitative results are necessary for computer simulation to be considered useful in theory construction – that the simulations must give concrete figures which can be compared against direct observations. This appears to claim that if concrete predictions cannot be made from it, then the simulation is not of much value.

This is hard to accept, and in many cases there are strong reasons for looking only for qualitative results. It must be recognised that models are abstractions of reality, with limited richness in possible interactions and accuracy. Further, most models will rely on single numerical parameters in place of factors which are not easy to accurately quantify – emotional state or the chance that someone might misinterpret some signal (see, e.g., Doran, 1996; Steels and Kaplan, 1998).

Most ALife also, at least implicitly, subsumes many ideas from work on complex systems – and accepting complex systems theory places limits on the predictability of such systems (e.g. Gleick, 1987). Complex systems often recognise that for different values of parameters, different classes of behaviour may exist. Thus qualitative predictions can be made where precise quantitative predictions may not be possible. The qualitative outcome may be predictable in many cases where it is not possible to put precise figures on the predictions.

So, quantitative results may not be feasible in many cases. Yet, where they can be obtained, they can certainly be compared more readily with collected data. Thorough statistical comparisons may be performed and the model tested against reality. Strongly correlated results increase confidence in the correctness of the model and theory; poorly correlated results might indicate that the theory is simply wrong.

3.8 ALife and the Evolution of Language

We previously highlighted the assertion of Taylor (1998) that the methodology used in any program of ALife research should be adapted according to the particular problem domain under investigation. Having reviewed methods and guidelines for the general application of ALife, it is now appropriate to add any special considerations required for using ALife in research into the evolution of language. Naturally, this depends on some understanding of the special problems, questions and methods that apply to the evolution of language.

Perhaps the single most important factor, as noted in Chapter 2, is that the evolution of language refers to two different processes: one biological and one cultural. Any ALife model should be aware of this, and clear in which elements of these different processes relate to the particular theories being explored. Work may focus on the biological evolution of linguistic ability in *Homo sapiens* or on the historical evolution of languages, or may involve both. As a result, we similarly divide this section on methodological considerations for computer modelling of the evolution of language, starting with concerns relating to the EoL.

3.8.1 ALife and the Biological Evolution of Language

Hurford (1992) presents six principles for evolutionary explanations of language. The application of these is not limited to computational modelling of the evolution of language, but could apply more generally to mathematical or other attempts at developing explanations for the evolution of language. These rules do, however, apply strictly to the biological EoL – thinking of language as an organ, with innovation by biological mutation and selection. In attempting to give an evolutionary explanation for the existence of any feature of language, Hurford argues that an investigator should consider the following factors:

1. Universality. The linguistic feature being investigated should be common to language in general.

2. Innateness. The feature should be innate, and evidence provided of this.
3. Contingency. A range of hypothetical alternatives should be presented and tested. In some cases only one or two alternatives may be envisaged. The more alternatives, the more impressive if the feature chosen is the only one that emerges.
4. Genetic Expression. Some relation of language feature to genes has to be made, and this should be stated.
5. Adaptive Value. The possibility space of alternatives should be related to fitness and or reproduction.
6. Demonstration. An argument should be presented that given language-gene mappings (4) and language-advantage mappings (5), the feature will either necessarily or probably emerge as the only survivor out of the possible alternatives (3). This can be presented as argument or simulation.

While these form useful guidelines, they are not without problems. The first point is clear and without problem – if the feature of language is not truly universal then it is unlikely to be innate (unless of course there is some evidence that possessing or lacking the feature is determined by heredity). But, as discussed in the previous chapter, determining the ‘innateness’ of a particular feature can be more difficult than establishing its universality. Further, the relation of an innate feature to genetic expression may be very hard to satisfactorily prove.

Depending on the features under investigation, the final principle may also prove difficult, where there are several possible outcomes. Given that chance has had a role to play in the EoL, it is possible that many features of language may not have been the features *most likely* to survive, but rather the ones that simply *happened* to survive (a point which Hurford acknowledges).

Despite these problems the principles are generally worthwhile, and emphasise the need to put effort in particular into showing that the feature under investigation is actually one that is a biological feature of language – an important prerequisite if we are to attempt to provide an evolutionary explanation for its presence. This is especially required to avoid proposing a biological solution to some culturally emergent aspect of language.

We also distinguished the evolution of language from the evolution of communication, and this is another distinction that we must ensure is embodied in ALife models of the EoL. Recognising that there are important differences between

language and other forms of communication, and that many special considerations apply to the evolution of language separate from more general questions on the evolution of communication, these must influence the design of ALife models.

A rhetorical example is in order here to illustrate this requirement. An attempt is made to demonstrate that language evolved due to a need to be able to warn others about possible attacks. A model is constructed, and results gathered which seem to support the hypothesis. However, such communication systems are also used by many species of animal (Hauser, 1996). Without additional work, or significant changes, there is no valid reason to assume that the results of the model are particularly relevant to the EoL.

One solution to help avoid this problem is to refer to the characteristic features of language (Section 2.2.1). Incorporating selected features, present in human language but absent in animal communication systems, can enforce the relevance of the model to the EoL. In doing so, it is not necessary to attempt to include all features. The fact that the majority of the features exist in some form, even if not all together, in a variety of animal communication systems should not prevent this approach from being useful. Instead, sets of features may be selected which together are only rarely, if at all, found together in non-linguistic forms of communication

Perhaps, in the previous rhetorical example, the author might seek to model a signalling system in which the repertoire of warnings is not fixed, but learned. Additionally the repertoire may not be of a fixed size, but may change.

Also extending the model, perhaps other ways of using signals could be incorporated, since deciding on a particular function that is served by language in the model is itself problematic. While language is known to serve many functions in human society, the problem may be caused not by ignoring other uses of language, but because of the problems in interpreting the effect of the single use modelled. Another example may be helpful.

A model is built in which successfully using language allows agents to somehow cooperate with one another. Such cooperation is rewarded by the model. A simulation is run and results gathered which show how language may evolve to support cooperation. With one explicit use of language built into the model, it may be that the space of evolutionary possibilities is quite small. Sooner or later, some agents will evolve the ‘linguistic ability’, and will start to cooperate. This is rewarded, and language evolution succeeds. With such a constrained evolutionary space and

rewards built in for the single anticipated and possible use of language, the model is simply showing the successful evolution of cooperation, not the evolution of language. It may actually be better to not include any explicit use of language at all. This problem is revisited later, in Section 5.7.4.

Some questions on how the evolution of the biological LAD might be modelled are discussed shortly, after turning our attention briefly from biological evolution altogether.

3.8.2 ALife and the Cultural Evolution of Languages

A different set of guidelines will apply to attempts to model the historical eol. Hurford's six principles will explicitly be irrelevant. What should also be inapplicable are the problems of confusing the evolution of language with the evolution of communication or cooperation – but this may not be the case.

Some ALife models, possibly using the 'meme' paradigm (Dawkins, 1976), may use hereditary signalling systems to represent language. Where the means of transmission incorporates learning from other agents there is no real objection to the use of memetic concepts. Where the translation from a genetic to a memetic model takes place without acknowledging the differences in methods of transmission, questions arise about whether cultural evolution is indeed being modelled.

The model developed in Arita and Koyama (1998), although not a memetic model, uses a hereditary signalling system although it has been developed to investigate the evolution of dialects in language. Unsurprisingly, the findings correspond more closely to the results that would be expected in investigations of the evolution of cooperation than the eol: common dialects allowing cooperation exist where resources are plentiful; non-compatible dialects (or even a complete lack of signalling) preventing cooperation exist where resources are scarce. Such findings obviously show how different environmental conditions can affect the evolution of cooperation in species – but the model fails to allow for non-cooperation between agents that share a common dialect. And evidence from studies of language diversity show the opposite trend from that described by Arita – scarce resources lead to lower, not greater, linguistic diversity (Nettle, 1999a).

So, an ALife model of the historical evolution of languages should avoid the genetic transmission of languages, and instead incorporate more realistic means of cultural transmission. An explicit learning mechanism is a better alternative – agents can

learn language from a selection of other agents instead of inheriting it from two parents. ALife models of the evolution of language should differ according to the particular aspects of language that are to be investigated, whether biological or cultural. But what of including both in the one model? While it is recommended that ALife models should be minimal, are there cases in which a single model should use both forms of evolution?

3.8.3 Coevolution of Language Ability and Languages

In the previous chapter the idea of coevolution was introduced, and any attempt to investigate this using an ALife model will require that the signalling schemes used in the model can evolve, as well as the innate linguistic ability. An example of this is found in (Briscoe, 2000b), which considers the coevolution of language and the LAD. Such models are quite complex, as many different interactions that affect the outcome are taking place. As well as the interactions between different agents as one part of the model evolves, there are the interactions between the coevolving elements of the model.

Additional complexity also introduces additional uncertainty into the results and the understanding of how they arise. To compensate for this, it is recommended in such models to test individual elements in isolation to verify their behaviour. The idea of unit testing an ALife model was mentioned earlier, in Section 3.4.4, but is extended here. As well as the testing of individual agents, the testing should also verify the outcome when evolution is only possible in one of the two coevolutionary systems.

For example, in a model like that described above, a test should be carried out on the learning of signalling schemes learned with a fixed, non-evolving, LAD.

3.9 Summary

There are a number of problems facing ALife researchers. The review given here highlights many of the problems, and gives a variety of solutions. While it is not possible to prescribe a single methodology for all ALife work, a high level of awareness of the problems facing ALife research is clearly desirable. Guidelines offer an easy way to promote awareness, and to *suggest* possible solutions. Undoubtedly, the sooner into a research project that a suitable method is adopted, the better. While the current state of methodological awareness in the ALife community is quite poor, interest in and awareness of methodological issues is growing. Some

problem areas still remain poorly covered in the literature – the academic pursuit of knowledge appears to be uncomfortable with software engineering. Problems theoretical and philosophical are more academically attractive than the mundane problems of ensuring code is correct.

In the descriptions of experiments and models that follow, details are given of some of the steps taken in verifying and validating the models and the principles outlined in this chapter are generally followed. There are, however, no rigorous software engineering development documents. No software quality process has been followed. In future years as models grow larger and more detailed there may be increased need for more rigorous approaches to software development in ALife. But for now precedence is rightly given to attempting to develop an emerging ALife scientific method, and to trying to show how a collection of ones and zeroes can be relevant to processes taking place outside the computer.

In the next chapter, the base model, round which the work of this thesis is developed, is presented and documented – according to many of the steps and methods prescribed in this chapter.

Chapter 4

An Artificial Life Model for Investigating the Evolution of Language

4.1 Introduction

In this chapter I develop the model that will be used in my computational investigations on the evolution of language. I describe the model, detail how it works and justify the decisions taken in implementing it. Some preliminary experiments are also detailed. The model as detailed here provides a base for further experiments on modelling the biological and cultural evolution of language, described in the following chapters.

4.2 Computational Models of Language Learning Populations

A number of recent computational models demonstrate the evolution of innate communication schemes (for example, Oliphant, 1996; Di Paolo, 1997a; Cangelosi and Parisi, 1998). Other models demonstrate the self-organisation of lexicons, grammars and sound systems in populations of language agents without evolution acting on the language agents themselves (e.g. Steels, 1996a; Batali, 1998; de Boer, 2000; Kirby, 2000). These models typically demonstrate the self-organisation of language features through a process of repeated interaction and learning.

Cangelosi et al (2002) describe a related model in which agents learn lexical terms for distinct visual inputs – square, ellipse, etc. This differs from the other models described here in that the lexicon is predetermined rather than emergent. This is a consequence of the model being used to examine the phenomenon of symbolic grounding (Harnad, 1990), not linguistic self-organisation.

Batali (1994) combines evolution and learning in an artificial neural-network, ANN, model in which recurrent neural networks attempt to learn context-free grammars in an investigation of innate language biases and critical periods. The language agents have a fixed structure, and a predetermined number of inputs and outputs. Evolution determines initial weight values for the networks, selecting appropriate values for the class of languages on which the population is trained. This model demonstrates how evolution can tune innate learning mechanisms towards certain grammars, once the mechanisms for language have developed.

A second model using ANN agents learning to produce and interpret signals, with observational rather than reinforcement based learning, is presented in (Oliphant, 1997). In Oliphant's model, ANN agents relate meanings to signals and vice-versa with winner-take-all competition on the produced vector, with meanings and signals represented by binary vectors with only one active value. Using a form of Hebbian learning (learning which increases the strength of a weight when the neurons which are connected by it both fire simultaneously), Oliphant shows successful negotiation to a common optimal language, with a different signal being used for each meaning. Another model where a community of ANNs negotiate a shared lexicon is presented by Hutchins and Hazelhurst (1995). The agents within this model are similar to the ones presented in this thesis, although their agents are more complicated with an additional layer (used only in learning). The authors limit their investigation to the development of a shared lexicon.

Similar, but not ANN-based, work has been presented, with a number of variations and enhancements by (Steels, 1996a; Steels, 1996b; Steels, 1996c; Steels and Vogt, 1997; Steels, 1998; Steels and Kaplan, 1998). In these works, agents attempt to learn symbol-meaning pairs through a negotiation process termed the 'naming game' – in which it is assumed that extra-linguistic means, such as pointing, enables agents to know the correct meaning of an utterance. While this does not capture the structural or generative nature of human language, these works have shown the emergence of shared lexicons in a number of distinct implementations.

Building a model where agents use simple grammars is certainly possible – (e.g. Elman, 1993) – but would introduce additional problems to be overcome. The first problem is simply that recurrent ANN, used for ANN grammar learning, have longer training periods than simple ANN. Using populations of ANN, this would have a significant overhead, and experiments would take significantly longer to perform. Using a simple signal-meaning mapping also allows for much simpler modelling of the evolution of language ability.

With a recurrent network, there is not a simple relationship between the number of nodes, or the connections between nodes, and the ability to learn a grammar. The networks themselves may consist of many layers and hundreds of nodes. How can the EoL be modelled with such networks? Batali (1994) achieves this, but in part only. By using a fixed ANN structure, initial weight values are evolved to find those best suited for language learning. However, given suitable training there is no reason

in principle why agents in the first generation should not be able to learn the languages eventually learned by those agents which have benefited from many generations of evolution. There is no evolution in the structure of the agents' ANN, only a tuning of their initial parameter settings.

In ignoring grammar, an important feature of language is lost but the details of the model are kept to a minimum and focussed only on properties and processes that are necessary for a minimal form of language.

An interesting approach would be one in which the ANN structure – the number of nodes and their interconnections – evolved. A number of approaches have been developed for evolving ANN of some complexity (for example, see Miller et al., 1989; Kitano, 1990; or Belew et al., 1991). These approaches all evolve ANN, train them to solve a certain task to determine fitness before producing a new generation of networks and beginning again. A difficulty in trying to use such models in investigations of the EoL is that, rather than evolving ANN to solve a single problem, the ANN need to *co-evolve* to solve a problem that itself varies according to the current population makeup. The nature of this problem is itself more complex.

4.2.1 Evolution of Language Ability

What the models described above do not show is the evolution of linguistic ability over time. The agents in later generations are not intrinsically 'more able' to learn language than those in the first. A slight exception is the work of Batali, but here it is simply the initial weights which are modified - the network structure, the capacity for language learning, is not.

In contrast, in Fyfe and Livingstone (1997) a population was modelled in which the capacity for language learning does evolve, with evolving network structure. In this work, the individuals first learn to identify stochastic sources in an environment and then learn a common language to communicate about the sources present in any given environment.

In this model, the language agents consist of three layers of nodes with two layers of interconnecting weights. The different layers of nodes store vectors representing the external environment, the signal generated in response and in the middle an internal representational state. This allows agents to use different internal representations from each other for the same meanings, without preventing them from reaching consensus on signals to represent particular environments.

Agents with fewer neurons in the hidden layer might be unable to accurately represent the different sources existing in the environment internally. Thus, they would be structurally less able to communicate effectively about their environment as would agents with larger hidden layers.

Experiments were performed with communities of agents with differing representational capabilities. Communicative success was seen to improve in populations with a common representational capability, and evolution towards homogenous representation capability was observed – but not necessarily towards better representation.

A weakness of this model is that the production of signals was compared to assess the success of communication, but the ability of agents to interpret signals was not tested. Thus, by producing the same word for the same environment two agents are assumed to be successfully communicating. No pressure to produce different signals for different environments, or to be able to decode signals, existed. While this was presented as a positive result, showing the emergence of synonyms, it is a consequence of the lack of pressure on agents to be able to interpret the signals and could be considered an artefact particular to this implementation.

The original aim of my research was to develop a new model in which the language ability could evolve over time, to allow investigation of factors surrounding the EoL, yet which overcame the problems of Fyfe and Livingstone (1997). In such a model, as the language ability of agents changed, so would the languages that they used. Thus, the model was to be one of language-physiology co-evolution.

4.3 Modelling Human Language

The model should capture some of the characteristic features of language (Section 2.2.1) if we are to claim that the results of running the computational model are relevant to an investigation of the evolution of language.

Ideally the model would include all relevant features of human language. Such a complicated model might, however, be very difficult to build and even more difficult to understand. As we saw in Chapter 3, models are abstractions of real world systems and the power they have in allowing us to improve our understanding of such systems comes from the ability to simplify, to remove complicated and possibly extraneous features, and to focus on selected aspects of a system. This applies equally to mathematical and to ALife models.

Thus, starting with a list of relevant features, we can select the minimum set that we feel will enable us to model the EoL. We saw in Section 2.2.1 that six key, text-book, features of human language are that language is:

- Communicative
- Arbitrarily symbolic
- Regularly structured and structured at multiple levels
- Generative and productive
- Dynamic
- Transmitted by learning rather than hereditary means

Of these features the first, that the signals need to be communicative, is clearly important. From an Information Theory perspective (Shannon and Weaver, 1949), a signal is communicative if it reduces uncertainty about the environment or some signified event. Thus, except where the same signal is produced for all possible sources, or where there is no causal relationship between source and signal, any system that generates signals from inputs will be communicative. However, an ideal communicative system removes *all* uncertainty. If the agents in our model are able to correctly interpret the majority of signals received then we will have established that the signals are highly communicative. (Although, see Di Paolo (1997b) for an account of the role of communication for coordination where there is no hidden information.)

Additionally, the agents in our model should acquire their arbitrary meaning-signal pairs through a learning process, ensuring that the signal scheme is arbitrarily symbolic and features ‘traditional transmission’. Such a system is extremely rare in natural communication systems other than human language. Few animals use signalling systems with discrete meanings and signals, and those that do appear to have innate rather than learned mappings (Hauser, 1996). Without at least these properties, a model may be more truly one of innate communication systems rather than of human language.

While it is possible to develop computer based models in which agents learn signals which are regularly structured (e.g. Batali, 1994), such models are more complex. As we will be attempting to develop a model in which the linguistic ability of agents evolves, such extra complexity could cause difficulties, and so the decision was made to focus on the already identified core features. As mentioned above, these features

are rarely found together outside of human language. Accordingly, they provide a reasonable set of features for building a valid model of the EoL.

Due to the abstractions and limitations inherent in any computer program, developing a model in which the language is generative and productive, allowing original and infinite use signals and combinations of signals is much more difficult. Within the bounds of a model in which populations of signallers exist, developing the models of the individual agents in order to allow the production and interpretation of a near-infinite set of signals is clearly not practicable.

The only remaining characteristic feature of language is that it is dynamic, changing over time. We investigate the dynamic nature of the signals used in our model in Chapter 6.

Finally, we will use the established ‘naming game’ (Steels, 1996a) interactions between agents for the purposes of language learning and testing.

4.4 An Artificial Neural-Network Based Language Agent

The individual language agents are implemented as simple ANN. The principal advantage of an ANN implementation is that it is relatively easy to generate individuals with differing network structures, representing differing innate linguistic abilities. A suitable learning rule will allow the development of language by individuals. As will be later seen, the model also features evolution in the communication schemes being used, akin to historical language change over many generations (Chapter 6).

It is easy for an ANN to learn uni-directional mappings, e.g. from some nominal ‘meaning’ to produce a ‘signal’. A learning rule was chosen which would allow ‘signals’ to be fed backwards to produce meanings, in a way similar to a Bi-directional Associative Memory (see Haykin, 1998). The fact that there is no pressure or, indeed, ability to make the signals model the meanings captures the arbitrary symbolic nature of language.

For the purpose of investigating the EoL, the previous three-layer model is not required; a simpler two-layer (internal state and signal) model can be used instead. While the agents no longer have differing internal representations, the function performed by the weights is the same – learning a mapping to generate common signals for common input states. Explicitly, in Fyfe and Livingstone (1997) the agents map an environment, e , to an internal state, i , to a signal, s :

$$\begin{aligned}
 s &= f_2(i) = f_2(f_1(e)) \\
 &= f(e)
 \end{aligned}$$

Simplifying this, it is possible to use a single layer of weights to perform a similar learning task. In Section 3.4.2 we saw that minimal models are to be preferred over more complex ones. The extra layer in the Fyfe and Livingstone (1997) model is more a distraction than an essential element of the model. In the following we assume that the function mapping from external environment to internal state is a given, and that agents have a common representation for different meanings – although as has been shown (Fyfe and Livingstone, 1997; Steels and Vogt, 1997), this is not required for ANN to negotiate a shared signalling scheme.

Having decided on the type of communication to be modelled in the simulated system, the next question is how can the success (or otherwise) of the evolution and emergence of ‘language’ be measured and evaluated.

4.5 Evaluating Communication in Artificial Populations

In the model I will develop, communication allows an agent possessing some internal state, or meaning, to send a signal to another agent. The receiver then maps the signal into an interpreted meaning. If the interpreted meaning is the same as the initial meaning, then communication can be said to have been successful.

One obvious measure of communicative success is to measure over a number of communicative episodes what proportion of signals sent are correctly interpreted. If all signals are correctly interpreted, then communication can be said to be perfect.

(Oliphant, 1997) defines three requirements which must be fulfilled for optimal communication, where every signal sent is correctly received. These are:

- Coordination. An individual’s meaning-to-signal generation must be coordinated with the inverse signal-to-meaning mapping.
- Distinctiveness. The signals used by an individual to represent each meaning must be distinct from one another.
- Consensus. All individuals in the population must have the same communication system. A caveat to this requirement is the assumption, made explicit by Oliphant, of equal numbers of signals and meanings.

The degree to which these requirements are met can be measured to evaluate the optimality of the communication system. Alternatively the simpler measure,

mentioned above, of success of interpretation can be used. To do this all that is required is to form pairs of agents where one generates a signal and the other interprets it. The percentage success can be measured over many pairings.

Information theory (Shannon and Weaver, 1949) provides further means for mathematical analysis of the communication schemes. The usefulness of the signals for transferring information between two agents is evaluated on the basis of mutual information in the meanings held by the two agents after signalling. Information theory also provides means to measure the information bearing capacity of the *channel*, as well as its quality (In Chapter 6 information theory is also used as a measure of linguistic consensus and diversity).

4.6 Modelling Language-Physiology Coevolution

In the model, agents will learn to map messages that are sent by other agents to internal states or ‘meanings’. By learning from each other a co-ordinated communication system is developed by a community of agents. The same ANN is used for the production of messages from each arbitrary meaning as for the reverse mapping.

In the next chapter I will add to the model by making the expressive capability of the language dependent on hereditary genes. The genes determine the number of language nodes possessed by agents, which determines the range of signals that can be produced. We can think of this as being a model of the coevolving languages used by agents, and the neural physiology that must exist to support the languages in use.

This will form the basis for investigating the evolution of language ability using the model. In the remainder of this chapter I will detail the basic model, without evolution, and present a number of experimental tests.

4.7 A Language Agent

With a requirement only that agents are able to learn signal-meaning mappings, a simple two-layer ANN architecture and learning algorithm is sufficient. If a standard linear ANN training algorithm (Haykin, 1998) were used for training individual agents, meaning-to-signal mappings could be successfully negotiated. But with no pressure on an agent being able to interpret the correct meaning of a presented signal the language could use a single word for more than one meaning. It would even be possible for a language to be negotiated that used only one word for all meanings (see

for example Fyfe and Livingstone (1997), a similar problem is also noted in Noble and Cliff (1996)). Using the criteria set out above, such a system would be far from optimal, fulfilling only the consensus criterion.

A learning algorithm which ensures the emergence of an optimal communication scheme - with a different signal shared by the whole population for each meaning state - is desired. The results of Oliphant suggest that optimal learning is performed by an algorithm using the transmission behaviour of the population to train language reception and the reception behaviour to train language production. Thus some kind of inverted learning algorithm is required. An inverse learning approach is presented in a number of generative models for ANN learning by Hinton and Ghahramani (1997). In these models, networks possess feedback generative weights and feedforward recognition weights, and the problem of recognition is posed as which hidden, or output, units could be responsible for generating the input pattern.

Using this approach, the learning algorithm should try to adapt the weights such that, for any given meaning-signal pair, the signal should produce the correct meaning when fed back through the ANN. We apply this principle to ANN with only a single set of weights, using the same weights for both the recognition and generative tasks. The operation of the learning algorithm is then as follows. A learner is presented with a meaning-signal pair. The signal is presented at the output layer of the learner and fed back to produce a *generated meaning*. This is compared to the original meaning, and any error is used to update the current weight values of the network (Figure 4.1).

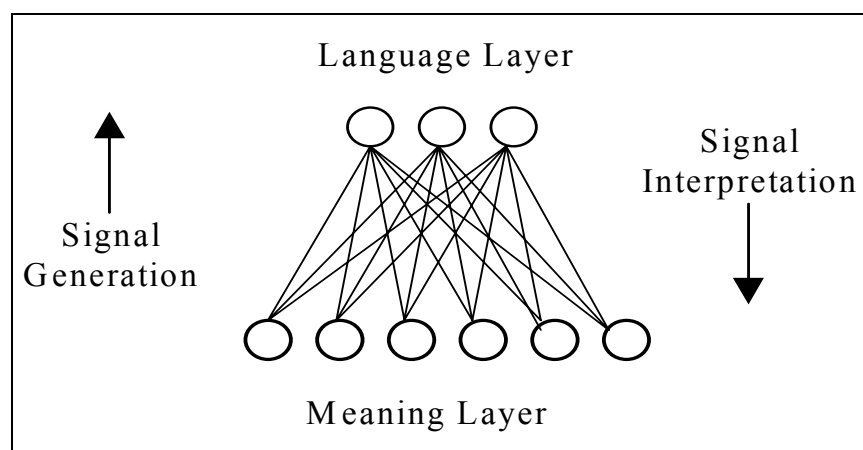


Figure 4.1. A language agent neural network. Recognition weights map language layer patterns to meanings, generative weights map meanings to patterns at the language layer.

The following sections present the equations for signal production and interpretation and for the learning algorithm.

4.7.1 Signal Production

An individual agent within the population is a single two-layer fully connected ANN with layers containing M and N nodes, respectively. ‘Meaning’ is modelled as a bipolar (± 1) vector of length M , which is presented at the inputs of a language agent, representing the agent’s internal state. Signals are similarly represented at the language layer as arbitrary bipolar vectors of length N .

A meaning vector can be fed forward through the ANN to determine an agent’s signal for that meaning, each output being thresholded to a bipolar value (± 1), as given below

$$y_j = \sum_{i=1}^M x_i w_{ij} \quad (4-1)$$

then $y'_j = 1$ if $y_j > 0$, $y'_j = -1$ if $y_j < 0$, $y'_j = \text{rnd}[1,-1]$ if $y_j = 0$

where the vector \mathbf{y} , $y'_1 \dots y'_N$, is the word generated for meaning vector \mathbf{x} . In all cases described, a sparse coding of the meaning is used with only one bit in the vector having the value $+1$, all others being -1 (However, the signal may be *any* arbitrary bipolar vector of length N).

4.7.2 Signal Interpretation

To interpret a signal vector, the signal can be fed back to generate a meaning vector. Competition can then be applied to set one, and only one, bit of the vector to $+1$, the remaining bits to -1 .

Thus, for N language neurons, there are 2^N possible signals or words in the language, and for M meaning neurons there are M possible meanings. Competition exists between neurons in the meaning layer, such that any signal fed back from the language layer only has one corresponding meaning. When a signal is fed-back during interpretation, it is likely that several meaning neurons will fire at different strengths, the competition allowing a single meaning to be chosen unambiguously.

For each neuron of the meaning (input) layer, equation 4-2 is applied to determine its activation value. Due to competition, the single neuron with the greatest activation value is set to $+1$, the remainder to -1 .

$$x_i' = \sum_{j=1}^N y_j w_{ij} \quad (4-2)$$

$$x_i' = 1 \text{ for } i = \arg \max_k x_k,$$

$$x_j' = -1 \text{ for } j \neq i$$

4.7.3 Learning

During learning, an agent will be presented with a meaning-signal pair. The signal will be presented at the output layer and fed-back to produce a generated meaning vector as described above. The error between the actual meaning, \mathbf{x} , and the generated meaning, \mathbf{x}' , is used for learning by the receiver agent, (4-3). This error is multiplied by a learning rate, α , to determine the correction to be applied to the weights, \mathbf{w} , connecting the layers.

$$\Delta w_{ij} = \alpha(x_i - x_i')y_j \quad (4-3)$$

This learning algorithm only updates weights when a word is misclassified. When a word is correctly classified the receiving agent performs no learning, as the error vector is composed of zeroes.

In the initial conditions, all weights have a zero value. Accordingly, signals will initially be random bipolar vectors as each output bit will be set to a random value (Equation 4-1).

4.7.4 A Note on Terminology

It is typical in ANN literature to describe an NN by the number of layers of *weights* it contains. Thus, a network with inputs fed through a layer of weights to neurons, then a second layer of weights to output neurons would be a 2-layer network. In this work, the ‘inputs’ can be applied to either side of the network, and are fed through a single layer of weights to neurons at the other side. Although there is only one layer of weights, we consider the network to have two layers of neurons, the two layers being the ‘meaning’ and ‘signal’ layers of neurons.

4.7.5 Representational Capacity

As noted, meaning vectors are sparse. So an ANN with M nodes in the meaning layer can learn M possible distinct ‘meanings’. Signal vectors are arbitrary, and for N signal layer nodes, 2^N possible signals can be learnt.

So a network should be able to learn eight meaning-signal pairs if it has eight meaning-layer nodes and just three language-layer nodes. This is not possible, however, as an additional bias node is required in the output layer, to be used in interpreting signals, as detailed below.

4.7.6 A Bias for Language Learning

This requirement for a bias bit can be explained. To simplify the explanation, consider an ANN with four input (meaning) and two output (signal) neurons. Such a network should be capable of learning four meaning-signal pairs.

For this case there are four possible two-bit signals, and these can be mapped onto two dimensions (Figure 4.2).




Figure 4.2. The dashed line through the origin fails to separate out one point from the other three. The solid line discriminates the top right point.

After feeding back any signal, one and only one bit of the meaning vector should be set to +1, the remaining three bits to -1. Two weights – the number of weights which feed back to each signal neuron – can code for a discrimination line of equation $ax+by = 0$. Such a line, through the origin, would be unable to separate off any one signal uniquely. By adding a bias bit to the signal, a discrimination line of equation $ax+by+c = 0$ can be derived. Such a line can discriminate one signal (Figure 4.2).

The above explanation can be extended to three dimensions to explain why three bits plus one bias bit are required in the output, signal, layer to allow the agents to learn eight pairs of meaning-signal mappings.

In summary, for M meaning nodes, the minimum number of signal neurons, N , including any bias, to be able to learn the bi-directional mapping is

$$N \geq (\log_2 M) + 1 \quad (4-4)$$

where M, N are integer values.

4.7.7 The Individual Language Agent – 1

The need for a bias bit can be confirmed by training an individual language agent on an arbitrary communication scheme. This is shown below in table 4.1. Each meaning vector has a -1 in all but one position, and it is the index of this position that is listed in the *meaning* column of the table. The signals have been arbitrarily selected as the binary numbers 0 through to 7.

Meaning	1	2	3	4	5	6	7	8
Signal	- - -	- - +	- + -	- + +	+ - -	+ - +	+ + -	+ + +

Table 4.1. An arbitrary communication scheme. Each of eight possible internal states is mapped to a different three-bit signal vector. (The bipolar values are represented as + and -).

The language agent is trained for four hundred rounds. In each round a randomly selected meaning-signal pair is presented once, and the error signal used to update the weights. The learning rate used is $\alpha = 0.2$, decreasing over the learning period by $d\alpha = 0.0005$ after each learning example.

Over a large number of runs, the agent continually fails to learn the mapping. Typical results are shown in Table 4.2.

Meaning Learned Scheme	1	2	3	4	5	6	7	8
1	- - -	- - -	- + -	- + -	+ - -	+ - -	+ + -	+ + -
2	- - -	- - -	- + -	- - -	+ - -	+ - -	+ + -	+ + -
3	+ - -	+ - +	+ + -	+ + +	+ - -	+ - +	+ + -	+ + +
4	+ - -	+ - +	+ - -	+ - +	+ - -	+ - +	+ + -	+ + +

Table 4.2. Results of four attempts to learn the communication scheme of Table 4.1. All attempts to learn the scheme fail.

4.7.8 The Individual Language Agent – 2

By adding a bias unit to the language layer, active during signal interpretation, it is possible for a language agent to learn any arbitrary mapping. The above experiment was repeated with the same parameter values, but with the addition of a bias node. Over a large number of runs, the language agent repeatedly learned the signal-meaning mapping, with no failures.

4.8 Experiment 1: Language Negotiation

The first proper test of the model is to determine whether homogenous populations of agents can successfully negotiate a communication scheme which they can use to share information about their internal states. To do so, an array is created and populated with language agents. Weights are initialised randomly. Over a number of rounds pairs of agents are selected, and communication and learning takes place. The algorithm is shown in Figure 4.3.

- | |
|--|
| <ol style="list-style-type: none"> 1. For t training rounds 2. pick random meaning 3. for each agent (picked in random order) 4. pick another agent to be signaller 5. generate training signal from signaller 6. train both agents on signal |
|--|

Figure 4.3. Population training algorithm

As detailed previously, with 8 possible meanings, a minimum of 3 language neurons, plus a bias, are required for a language to be able to convey all possible meanings.

4.8.1 Emergence of a Language Bias

Without setting a bias, it is still possible for populations to negotiate a language which allows the successful communication of all meanings – but only where the number of neurons in the language layer satisfies equation (4-4). E.g., for three possible signals three nodes are required at the signal layer.

Thus, successful language learning is possible for a system with eight possible meanings where the agents' language layers have no bias but at least 4 neurons. However, in all successful experiments with 4 language neurons, one of the signal bits becomes set to the same value (+ or - 1) in all signals. In effect, although there is no explicit bias, one emerges from the interactions of the agents as they attempt to learn to communicate with one another. This is shown in the negotiated languages of Table 4.3.

In Table 4.3(a), the signalling scheme shown is for one agent in a population of agents which *all* use the second bit of their signals as a bias – in all signals, after negotiation, this bit is set to -1 . A similar result occurred in the population from which the signalling scheme in Table 4.3(b) is drawn, only this time the fourth signal bit is always set to $+1$.

a) Word 1: 1 -1 -1 -1	b) Word 1: -1 1 -1 1
Word 2: -1 -1 -1 -1	Word 2: 1 -1 -1 1
Word 3: 1 -1 1 -1	Word 3: -1 1 1 1
Word 4: 1 -1 -1 1	Word 4: 1 1 -1 1
Word 5: -1 -1 1 -1	Word 5: 1 -1 1 1
Word 6: -1 -1 1 1	Word 6: -1 -1 -1 1
Word 7: -1 -1 -1 1	Word 7: -1 -1 1 1
Word 8: 1 -1 1 1	Word 8: 1 1 1 1

Table 4.3. Two negotiated languages. In (a) the second bit is uniformly set to -1, in (b) the final bit is set to 1.

Further evidence of the emergent bias comes from examining the internal weights of a language agent drawn from a population in which an optimal communication scheme has been learnt. The weights corresponding to the bias bit are large, on the order of three times the size of other weights. All other weights are the same order of magnitude as each other. The set of converged weights for one language agent are presented in Table 4.4.

Meaning, M	W_{i1}	W_{i2}	W_{i3}	W_{i4}
1	0.233852	0.749355	-0.25996	-0.24527
2	-0.23764	0.787846	-0.27316	-0.2321
3	0.299894	0.740592	0.269213	-0.23706
4	0.241851	0.717312	-0.2254	0.250896
5	-0.2961	0.722206	0.263901	-0.28557
6	-0.22819	0.744461	0.25074	0.291194
7	-0.23806	0.745487	-0.24148	0.226476
8	0.224402	0.792741	0.216147	0.231434

Table 4.4. The weights for a NN in a population which has negotiated that the second language bit act as a bias, with value -1. The language of this particular agent can be seen in Table 4.3(a)

4.8.2 Success of Language Negotiation

To evaluate language negotiation, populations were initialised and test runs completed for varying values of N and varying numbers of training rounds. All other parameters were kept constant – except for $d\alpha$, which is derived from the values of α and t , such that α decrements to zero over the training period. The experimental parameters are shown below in Table 4.5. Additionally, all agent weights are initially set to zero. Accordingly, the initial set of signals produced will be random. The populations of agents are then given either 100 or 250 rounds of language negotiation to form signalling schemes.

Fixed Parameters		Variable Parameters	
Learning Rate, α	0.2	N	0 to 6
Population	120	Training rounds, t	100 or 250
M	7	$d\alpha(\alpha/t)$	0.0005 or 0.0002

Table 4.5. Experimental parameters

For each parameter set, 10 runs were performed. Within each run success was measured as the percentage of signals correctly interpreted by the agents. Testing was performed, after training had completed, using the algorithm presented in Figure 4.4. During this testing no learning is performed by the agents.

1. For each agent, A, in population; %success = 0
2. For test = 1 to 100
3. Select random agent B, such that $A \neq B$
3. Pick random meaning, M
4. Present M to B to produce signal, S
5. Present S to A to produce generated meaning, M_G
6. If $M_G = M$, %success = %success + 1
7. Next test
8. Average success scores of all agents

Figure 4.4. Algorithm for evaluating language negotiation success

The percentage successes for different combinations of parameters are shown in Figure 4.5. The results show the standard deviations across the ten simulation runs, as well as the means, for each parameter set. The result obtained for success of communication for populations where the agents have no units in the language layer reflects the chance of a randomly guessing the meaning.

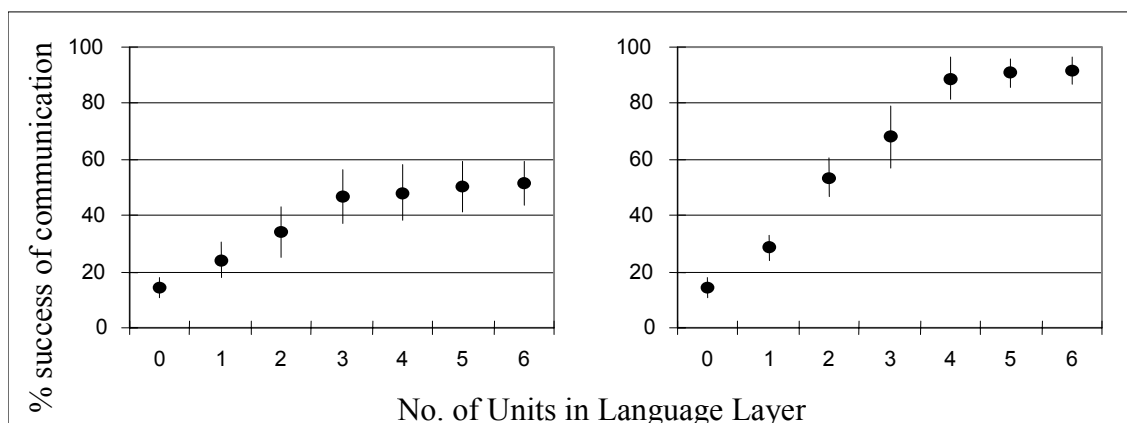


Figure 4.5. Success of language negotiation, with variable numbers of units in the language layer. Average success (with standard deviation) in interpretation shown. Left: 100 rounds of language negotiation. Right: 250 rounds.

4.8.3 Interpreting the Results

With no language neurons, the performance of the agents is at chance level - on average agents guess the correct meaning one in seven times. As language capability increases, so does fitness showing that agents are successfully sharing information about their internal state. With more neurons, average success increases and with two or more language neurons increased training time increases fitness. If training time is extended sufficiently, populations with four or more neurons will consistently negotiate a language capable of allowing information to be shared, with a very high communicative success rate.

The results show that using our learning algorithm homogenous populations are able to negotiate a useful language, even where the capability for an optimal language does not exist. As the language capability increases, the success rate of communication increases. With four or more language units, the expressive power of the language exceeds the communication requirements of the environment. This has the potential for allowing synonymy, where multiple signals for one meaning may be recognised correctly.

Without at least as many language nodes as required by equation (4-4), four nodes in the experimental setup described above, no population was able to negotiate an perfect signalling scheme. Given the finite number of training rounds, not all populations in which the agents satisfied equation (4-4) negotiated optimal communication schemes. But success at interpreting signals from other agents in the population is seen to improve as the number of neurons in the language layer increases.

External confirmation that the agents and learning rule used in this model are able to negotiate an optimal communication system is made in a recent work, which categorises the different possible learning rules (Smith, 2002). According to this, the learning rule used here has biases for allowing agents to learn and maintain existing optimal communication systems, and to create such systems where they do not already exist.

4.9 Experiment 2: Spatially Arranged Populations

In Chapter 2, it was noted that the EoL and the evolution of communication are related to the evolution of co-operation itself and this is explored more in the

following chapter, where we investigate the possible effect of spatial selection on the EoL. Rather than add spatial selection and enable evolution of agents simultaneously, it is worthwhile repeating the previous experiment to see if spatial constraints on agent communication have any impact on the results we obtain.

In the previous experiment, when two agents were to be selected to communicate for learning or testing, any two agents could be selected. In any and every pairing, every agent had a uniform chance of selection. Here the same experiment is repeated, but with a spatial organisation imposed on the population of agents. This will allow us to observe any differences before making any further changes to the model. All parameters have the same values as previously, and the language is still negotiated over a number of training rounds in a single homogenous population.

A simple spatial arrangement is created by organizing the agents in a ring, and limiting all communication to within a neighbourhood of the currently chosen agent. So, for example, a teacher will be picked randomly according to the location of the current learner. The area of the neighbourhood is defined by a normal distribution centred on the learner. In all experiments described the standard deviation used is 0.6, placing a strong preference on immediate and very close neighbours. A ring arrangement is not necessary – a linear arrangement could be used instead – but means that all agents have the same number of neighbours and have the same chance of selection.




Figure 4.6. A normal distribution curve determines the likelihood of a neighbour being selected as a partner in for learning or evaluating communicative success. An individual cannot learn or otherwise communicate with itself, and so the centre of the curve is ‘zeroed’. The standard deviation used is 0.6, effectively limiting communication to close neighbours.

With neighbourhood limited communication the results in Figure 4.7 were produced for 100 and 250 training rounds for populations of 120 agents, again each result averaged over 10 runs. Success is measured as detailed in Figure 4.4. The same

neighbourhoods used in language negotiation are used for communication during fitness evaluation.




Figure 4.7. Average success of signal interpretation, and standard deviation, with homogenous populations of 0-6 language neurons measured over ten experiment runs each case. With 100 (left) and 250 (right) negotiation rounds.

Comparing both sets of results, there appear to be a few small differences, which warrant some further investigation. Using spatial constraints on agent interaction appears to increase the variance in the success of language negotiation – the standard deviations shown in Figure 4.7 appear to be consistently larger than those in Figure 4.5. Additionally, the average success rates show some improvement when spatial constraints are present, particularly where the number of training rounds is low.

Why should the variance in success be greater, and why should the success rate of communication have improved? If the communication is limited to nearby agents, it may be the case that language negotiation is more successful amongst some agents than others in a single population. This could occur if different signalling schemes are negotiated amongst different subgroups of a population – a dialect-like effect. This could also explain an increase in success, as agents only have to negotiate a common signalling scheme with a subset of the population, rather than across the whole population. Communication between adjacent agents would have a higher success rate than communication between more distant agents. This could cause both the higher success rate, and the increased variance. Below we describe a further run with spatial limits enabled, and explore the signalling schemes of individual agents to determine if dialects are indeed formed.

4.9.1 Emergence of Dialects

It was noted previously that where agents possess a greater number of language layer neurons than required, it may be possible for them to interpret more than one signal

as having the same meaning. Communication may still be otherwise optimal, potentially with complete success in signal interpretation. If, in the last experiment with a spatially arranged population, dialects did emerge, this can be put to the test by examining the signals used by different agents, and measuring the success of agents at interpreting different signals.

A homogenous language population with four language units, plus bias, is trained for 5,000 rounds. This allows a very high degree of coordination amongst the agent languages but, due to the local communication, does not negotiate a common language over the whole population. It is observed that large neighbourhoods negotiate a common signal for a given meaning, but distant agents may have significant differences in communication schemes used. At the boundaries between neighbourhoods, agents may exist which interpret signals from different schemes correctly. This is shown in Table 4.6. The three agents included each attained a maximal fitness score, interpreting all signals correctly. This demonstrates that a degree of ‘bi-lingualism’ or even ‘multi-lingualism’ is possible in the agent communication schemes. The additional language layer nodes provide some redundancy in the representation capacity of the agents.

Meaning 0						Meaning 4					
Agent	4	-1	1	-1	-1	Agent	4	1	-1	-1	-1
Agent	5	-1	1	-1	1	Agent	5	1	1	-1	-1
Agent	6	-1	1	-1	-1	Agent	6	1	-1	-1	-1
Meaning 1						Meaning 5					
Agent	4	1	1	1	1	Agent	4	1	-1	-1	1
Agent	5	1	-1	1	1	Agent	5	1	-1	-1	1
Agent	6	1	-1	1	1	Agent	6	1	-1	-1	1
Meaning 2						Meaning 6					
Agent	4	1	1	1	-1	Agent	4	-1	-1	1	1
Agent	5	1	1	1	-1	Agent	5	-1	-1	1	1
Agent	6	1	1	1	-1	Agent	6	-1	-1	1	1
Meaning 3											
Agent	4	-1	-1	1	-1						
Agent	5	-1	-1	1	-1						
Agent	6	-1	-1	1	-1						

Table 4.6. The signals used by three adjacent agents for seven environmental states. Each meaning is represented by a bipolar (+ and - 1 values) signal vector, e.g. to indicate meaning 0, agent 5 would send the signal (-1,1,-1,1). All three scored maximum fitness, interpreting all signals correctly, despite differing communication schemes. Bias signal is not shown.

4.10 Conclusion

In this chapter I have developed a basic model for exploring the evolution of language and described some preliminary experiments. These have shown that the

agents in the model are able to negotiate a shared communication scheme, a simple ‘language’. This replicates work done on demonstrating self-organising languages and signalling schemes, as demonstrated by (MacLennan, 1991; Hutchins and Hazelhurst, 1995; Steels, 1996c) amongst others.

From Figures 4.6 and 4.8, another result is apparent, however. As the language capability increases, so does the communicative success rate of the agents. This is to be expected. But, as the language capability increases beyond what is required to express all of the available meanings, communicative success continues to improve. There would appear to be some benefit in having a linguistic ability somewhat greater than strictly necessary. To rephrase, a linguistic ability with some degree of *redundancy* has an advantage over one that is merely sufficient. This result, and some of its implications are reviewed as part of the next chapter.

The emergence of dialects, noted in the second experiment, is investigated more thoroughly in Chapter 6.

Chapter 5 Modelling the Biological Evolution of Language

In the previous chapter, we described the basic model which will be used as the basis of our experiments, and saw how the language agents are able to negotiate shared communication schemes.

In this chapter we adapt the model in order to conduct an experiment investigating the biological evolution of language ability. Changes to the model presented in the previous chapter are detailed, and experiments run and results presented. We conclude the chapter by critically reviewing the contribution that the experimental results and analysis has on existing debate on the Evolution of Language.

Before proceeding with the experiment, we begin by highlighting some of the key points in the debate, and show that previous work on modelling the emergence and evolution of language has not been sufficiently targeted at answering some of these points.

5.1 The Evolution and Emergence of Language

Many of the relevant aspects pertaining to the EoL have been reviewed in Section 2.4, and Chapter 4 contains brief descriptions of various attempts to build models of the emergence and evolution of language. It is worthwhile reiterating several points, to emphasise the goals of the following work.

It is clear that human physiology has been adapted in a number of ways which has enabled the learning and use of language, and that there are costs associated with a number of these changes. It is this evolution of physiology, which supports language, that we wish to model. As has been noted, few of the current models of the evolution of language incorporate this, demonstrating instead how languages may evolve in populations of capable language learners. Aside from the already noted exception of Batali (1994; 1998), those models in which individuals evolve as well as their communication schemes tend to demonstrate the evolution of innate, rather than learned, communication schemes.

Such models commonly relate the evolution of signalling or language to a single specific task or adaptive function. These are more akin to animal communication systems than to language. As well as demonstrating a model of the evolution of language, we will use our model to investigate some aspects of the evolution of language, particularly the influence that spatial constraints on the population have on the evolution of linguistic ability. We also see how some design decisions when implementing the model can have a significant impact on the observed results.

5.2 Modelling the Biological Evolution of Language Ability

One problem to be explicitly addressed in our model, but not in the other models already cited, is that communication and learning of language has to be possible with heterogeneous populations. Two agents of differing innate abilities for language learning and use must be able to learn to communicate with each other despite their differing abilities. This requirement arises as a result of the inevitable heterogeneity which must occur in a population as a particular trait evolves.

Once some mutation occurs to initially provide the trait in a few individuals, many generations may pass before it is shared by all. In the case of language, if individuals with more developed linguistic abilities are not able to form a common language with their less well developed neighbours then their abilities will not confer the expected fitness benefits. Without modelling the coevolution of physiology and language, it is not necessary to resolve the problems of how to represent evolving language ability, and of how to ensure that heterogeneous agents can negotiate communication schemes. Here they must be tackled and explicit solutions to these problems are required.

5.2.1 Heterogeneous Language Abilities

This can be captured using ANN based agents derived from those presented in the last chapter. The simple change required is to make the number of neurons in the language layer a hereditary trait, subject to change and evolution. Assuming all agents in the population have at least one neuron in the language layer, subsequent generations may feature heterogeneous mixes of individuals with one, two or more neurons, Figure 5.1.

All agents have the same, fixed, number of neurons in their meaning layers. While this makes modelling easier, it is acknowledged that this is not necessarily true in the real world. For a synthetic example, Steels and Vogt (1997) have also shown that shared internal representations are not required for successful communication about external objects common to multiple observers.




Figure 5.1. Over a number of generations the number of neurons in the language layer evolves. Even within a single generation, different agents may have different structures, and will have to be able to communicate despite the differences.

The following examples demonstrate how communication is possible between the three agents of Figure 5.1, assuming them to be members of a single population of agents.

With three nodes in the meaning layer, there are three possible meanings. With one, two and three nodes in the language layer they can produce two, four and eight distinct signals, respectively (4.7.2). A bias node is present in the language layer in all three cases but is not shown. For the three meanings, Table 5.1 shows signals that might be produced by the three agents.

Meaning	Signal Produced		
	(A)	(B)	(C)
1 (+ - -)	+	+ -	+ - -
2 (- + -)	+	+ +	+ + +
3 (- - +)	-	- +	- + -

Table 5.1. Signals which could be produced by the three language agents shown in Figure 5.1. The bipolar vectors are shown here simply as strings of + and -.

If a signal were sent from agent C to agent A, A would receive only the first bit of the three bit signal as it only has one node in the language layer. So, according to table 5.1, the communicative episode might unfold as shown in example (5-1), below:

Meaning 2 randomly selected
 Agent C generates signal vector (+ + +)
 Agent A receives signal (+) (5-1)
 Signal interpreted as meaning 1 or 2 (depending on activation values)

Conversely, if the signal received is of a shorter length than the agent is capable of, then it is padded with zeroes. A zero value presented at a language layer node will have no effect on signal interpretation when fed back – as can be seen from equation (4-2), a zero bit in the signal, y , has no influence on the interpreted meaning, x . Thus, only non-zero (+ or -

1) bits in the signal can have an effect on the interpreted meaning. Example (5-2) demonstrates this:

Meaning 2 randomly selected
 Agent A generates signal vector (+)
 Agent C receives signal (+ 0 0) (5-2)
 Signal probably interpreted as meaning 1 or 2 (depending on
 activation values)

As demonstrated above, using the agent architecture as described allows communication in heterogeneous populations. Such communication might not always be successful, but could improve the likelihood of successful communication significantly above chance. A further modification further improves the chances of successful communication.

5.2.2 Comprehension Leads Production

It has been argued (Burling, 1998) that comprehension leads production – that the ability to understand or interpret signals leads the ability to produce them. This is seen in the ability of many animals to understand commands given to them, in ape language learning and in human speech acquisition. Accepting Burling’s argument, it is desirable to also capture this feature in the model. This is also done quite simply.

All agents have identical structures, but hereditarily determined language production ability. This is implemented by limiting the number of neurons in the language layer which can be active during language production, any inactive neurons producing a zero value, Figure 5.2. All neurons are active for signal interpretation.

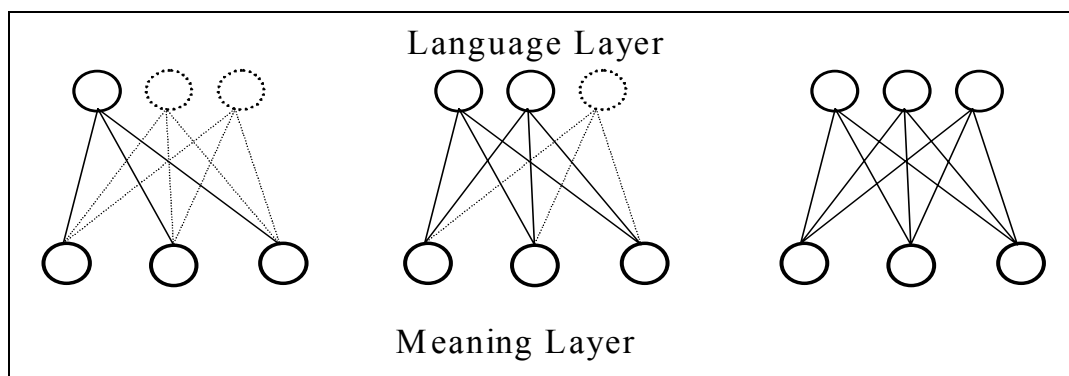


Figure 5.2. In the comprehension-leads-production model, all language agents have the same structure, but only a limited number of nodes may be active for signal production.

By implementing the changes described above, all agents are potentially able to learn to interpret all signals, including those which they are unable to produce. This is demonstrated in the following two examples.

If a signal were sent from agent C to agent A, A would receive all of the three bit signal as all nodes are active for interpretation. This communicative episode might unfold as shown in example (5-3):

Meaning 2 randomly selected
 Agent C generates signal vector (+ + +)
 Agent A receives signal (+ + +)
 Signal interpreted as meaning 2

(5-3)

During production, the signal is padded out with zeroes. Again, zero values presented at a language layer node will have no effect on signal interpretation when fed back, and so only the bipolar values have an affect the interpreted meaning. Example (5-4) demonstrates this:

Meaning 2 randomly selected
 Agent A generates signal vector (+ 0 0)
 Agent C receives signal (+ 0 0)
 Signal interpreted as meaning 1 or 2 (depending on activation values)

(5-4)

In this case the result is similar to the previous version (example 5-2), but where comprehension leads production poor signallers may benefit more from good signals produced by others than they would otherwise (comparing examples 5-3 and 5-1).

The experiments which follow are performed using this full interpretation/limited production design.

5.2.3 Population Generations and Replacement

The population training algorithm shown in Figure 4.3 will again be used. In these evolutionary models, training and fitness evaluation is performed for one generation. A ‘child’ generation is created, then this replaces the parent population. Training and fitness evaluation are again carried out, and a further child generation is performed. This process is repeated many times over, during which time the evolution of the language capability of the agents can be viewed.

The next few sections detail the design decisions, and parameter settings relevant to the genetic representation and agent selection and reproduction.

5.2.3.1 Genetic Representation and Reproduction

The standard genetic algorithm, GA, (Holland, 1975; Mitchell, 1996) is perhaps the most common technique used in modelling evolution. We use a GA in this model to allow the language capability to evolve over a number of generations.

The only hereditary trait to be modelled is the number of language neurons active during signal production. In the experiments performed here, the number of meanings is arbitrarily set to eight. As an agent with M language neurons, plus bias, can learn unique signals for 2^M meanings, the representation used should allow for at least zero to three language neurons.

In the standard GA, hereditary traits are represented by binary strings. Using three bit gene strings as the chosen representation, agents may have from zero to seven production-active language layer nodes. The standard binary representation of these numbers is as shown in table 5.2.

0: 0 0 0	1: 0 0 1	2: 0 1 0	3: 0 1 1
4: 1 0 0	5: 1 0 1	6: 1 1 0	7: 1 1 1

Table 5.2. The standard binary representation of the numbers zero to seven.

During reproduction, two parents' bit strings are selected from which two child bit strings will be produced. The standard operators for producing the child gene strings are *crossover* and *mutation*. As a result of crossover, each child string is part produced from a portion of one parent string and the remainder is taken from the other parent. Mutation may cause individual bits of the gene string to 'flip', from 0 to 1 or vice-versa.

The use of the standard binary representation in the GA has been criticised, however. For example, it is not always possible for a value to be incremented or decremented by one with a single bit-flip. To increment from the value 3 to 4, as shown above, all three binary bits must change. Thus, using the standard binary representation in the GA, many genetic changes may be required to produce a small change in the phenotypic form.

One alternative is to use a Gray coded (Gray, 1953) binary representation to store values in the genome, and this has been recommended for use with the GA by a number of authors (for example, (Belew et al., 1991) and (Caruna and Schaffer, 1988)). The Gray codes for the values 0 to 7 are shown below in Table 5.3.

0: 0 0 0	1: 0 0 1	2: 0 1 1	3: 0 1 0
4: 1 1 0	5: 1 1 1	6: 1 0 1	7: 1 0 0

Table 5.3. Gray coded binary representation of the numbers zero to seven.

Using Gray codes, only one bit change is ever required to increment or decrement a value by one. With reference to the GA, this confers the possible advantage of having similar phenotypic forms possess similar genetic representations. For modelling evolution, as opposed to the more general use of GAs in optimisation, are incremental changes particularly to be sought after?

In nature, relatively minor genetic changes *can* result in significantly different phenotypic forms. In general, however, individuals with similar genes may be expected to have more similar phenotypes than those with dissimilar genes.

If we wish our model to capture this feature of the natural world then a genetic representation which captures the adjacency feature provided by Gray codes is to be preferred over the standard genetic representation.

Gray codes may not be the best alternative. Although it possesses the benefit of adjacency, it is possible for even more dramatic discontinuities in evolution than are possible with the standard representation. Only one mutation is required to jump from the smallest possible value to the largest in a Gray coded gene string, from 0 to 7 in this case. Using the standard representation, the largest change a single mutation could cause would be an increase of 4, for example from 0 to 4 (Table 5.2).

The best option may be to dispense with a binary representation altogether. For example, Mühlenbein and Schlierkamp-Voosen (1993) detail various crossover and mutation operators for gene strings composed of real rather than binary values. This use of real values, rather than binary bits, in a genetic string has been used for many years in attempting to solve numerical optimisation problems under the name Evolution Strategies, ESs, an approach developed in the early seventies by Rechenberg (Mitchell, 1999).

However, our model uses discrete rather than continuous values to determine the evolved linguistic ability – with integer numbers representing the number of neurons active in signal production. Accordingly we can use integer valued rather than real valued gene strings to represent the innate linguistic ability of each of the agents. This will provide the adjacency benefit of Gray codes but without the possibility of single mutations causing such large changes.

Thus we have two different possible genetic representations for language ability – the standard binary representation, and an integer representation. In our first experiments we will see the effect of choosing one of these representations over the other, but first the mechanisms of selection, mating and population replacement common to both versions will be detailed.

5.2.3.2 Selection, Mating and Replacement

Other model details requiring elaboration surround the processes of selection and mating. The populations in all experiments described in this chapter are spatially arranged as described in Section 4.9, with the standard deviations used noted in each case. The

replacement and mating algorithm has been selected such that the spatial relationships between agents are maintained not just for learning and fitness evaluation, but also for mating and the placement of child agents into the succeeding population.

The fitness measure used is the percentage communicative success measure described in Section 4.8.2. Every agent has a fitness value calculated. Agents are then selected for reproduction, with each agent having a chance of selection equal to the agent's fitness score over the total fitness score for the whole population. The mating and replacement algorithm is as shown in Table 5.4.

For population size P , repeat $P/2$ times:

- Select parent, p_1 , randomly according to fitness
- Select parent, p_2 , randomly according to neighbourhood around p_1
- Mate p_1 and p_2 , producing children c_1 and c_2
- Place c_1 and c_2 in child population

Replace parent population with child population

Table 5.4. The mating and population replacement algorithm.

An agent may be selected as the first partner for mating more than once. The distance measure used for selecting a second mating partner is in all cases the same function as used for selecting a partner for signal learning or fitness measurement.

The replacement algorithm, for adding child agents to the new population, works a little differently.

Recalling the linear arrangement, if the first parent is 23rd in the line in its generation, the replacement algorithm will attempt to place the children in the 23rd and 24th positions in the child population. If either or both of these slots are already occupied, then the replacement algorithm iterates incrementally along the line until empty slots are found for both children. This replacement algorithm serves to reinforce and maintain the spatial relationships within the model.

Note that there is no other interaction between generations in this model. In particular, the language negotiation process starts afresh in each generation, with no input from past generations. Over time, it is this ability to negotiate language that is being evolved.

5.2.3.3 Crossover and Mutation

Reproduction in GAs utilises two key operations: *crossover* and *mutation*. The workings of these operators depends on the method of representation, as well as on design decisions (for example, see Mühlenbein and Schlierkamp-Voosen, 1993).

During crossover, the two parent gene strings are copied, with a portion of the gene string of each parent going to each of the offspring. During mutation additional random changes to the child strings may occur.

As the precise details of implementation depend on the chosen genetic representation, these are expanded upon later.

5.3 Modelling Language-Physiology Coevolution, Part I

The first experiments are to be simple demonstrations of the co-evolution of language and physiology, initially using standard binary codes in the GA. The extent to which agents are ‘physiologically’ adapted for communication is revealed by the number of active language production neurons each possesses, the fitness of a group of agents revealing how well the learned communication schemes have been negotiated for the sharing of information.

Over the course of many generations the language capability of the agents will evolve, and with it the signal repertoires available to them. The ability to successfully negotiate useful communication schemes in each generation being the driving force for the evolution of the ANN.

5.3.1 Crossover and Mutation

Using the standard binary representation, the gene string of each agent will be just three bits long. This is sufficient to store any integer value in the range $[0,7]$, and with 7 language neurons an agent can potentially learn 2^7 (128) different signal-meaning pairs. Where there will only be 8 possible meanings, as here, this should be quite ample.

The crossover and mutation operators used are quite standard implementations for use with the GA.

For crossover, assume two parent strings A and B. These will mate to produce two child strings C and D.

$$\mathbf{A} = \{a_1, a_2, a_3\}, \mathbf{B} = \{b_1, b_2, b_3\} \quad (5.5)$$

To produce C, crossover is performed such that, for each bit in the child string:

$$c_i = \{a_i\} \text{ or } \{b_i\} \quad (5.6)$$

Where a_i is chosen if $i \leq N$, else b_i is chosen. N , the crossover point, is a random integer chosen uniformly from the set $[1,2,3]$. The second child is a complement of the first in drawing each bit from the other parent:

$$d_i = \{b_i\} \text{ or } \{a_i\} \quad (5.7)$$

Where b_i is chosen if $i \leq N$, else a_i is chosen.

Mutation is applied after crossover has generated the two child strings. The *mutation rate* gives the chance of a mutation occurring at each position in the string. Each mutation that occurs inverts the binary value at that position.

5.3.2 Parameters and Settings

Table 5.5 lists the parameters and settings for the first coevolution test.

Parameters for coevolution (using standard binary representation)			
Learning Rate, α	0.2	Mutation rate	0.005
Population	100	Training rounds, t	200
M	8	$d\alpha (\alpha / t)$	0.001
Standard Deviation	1.0		

Table 5.5 Parameters and settings

The initial population of agents all have $N = 1$ (only one language layer node active for signal production). All agents have $M = 8$ nodes in the meaning layer. The mutation rate was selected so as to produce one or two mutations per generation. This might be considered a low rate of mutation were the GA being used for some optimisation problem, but here we wish to limit the number of significant mutations which occur in each generation of what is quite a small population. The standard deviation determines an agent's neighbours, as previously described. This gives chance greater than 80% of selecting an immediate neighbour, with over 98% of selections being within a small neighbourhood or two neighbours on either side.

5.3.3 Discontinuous Evolution of Signalling Ability

With the model as described, a typical run using the parameters shown in Table 5.5 produced the results shown in Figure 5.3. Displayed in Figure 5.3 are the averages of agents' fitness scores, and of agents' number of language neurons active in signal production.

The first demonstrates the effect of the second. Although an agent does not need active language production neurons to score highly on fitness, fitness being evaluated on the basis of how well received signals are understood, it does require that its neighbours have them.

In this figure the gradual evolution of language ability in the population is apparent. In a test run, over the course of 150 generations, the average number of active signal production neurons rises from 1 to 6, and the average success at interpreting signals from 20 to 80 percent (Figure 5.3).




Figure 5.3 Coevolution of language and physiology. As the communication capability increases (number of nodes), so does the success rate (fitness).

The graph shows a sudden jump in the linguistic ability and in the fitness of agents quite early on. This appears to be a significant discontinuity. Rather than a slow gradual increase in linguistic ability, it seems that very small number of mutations have led to major changes to the fitness of those affected individuals, and/or their partners in language negotiations and evaluation. This has resulted in a small number of individuals having much higher fitness, allowing the better adapted genes to rapidly take over the population. The graph is re-plotted for the first 40 generations only, in Figure 5.4, to allow a closer inspection.

In just ten generations (the 15th to the 25th) the average number of active language layer nodes rises from just over one to almost 5. With only one or two mutations to be expected per generation, the impression that relatively few mutations have produced rapid and major evolutionary development is confirmed.




Figure 5.4 Detail from beginning of Figure 5.3

However this is only the result of one experimental run. The runs themselves are highly random with many stochastic elements – which individuals get picked, which meanings they attempt to communicate and what signals they initially send before training.

Consequently, a large number of tests are required before the results can be considered at all robust or typical. Figure 5.5 shows the results of a further nine runs using the same experimental setup as in Table 5.5. It is apparent that, despite the individual differences, the overall result of significant improvement in communicative ability is consistent. In all cases the improvement is relatively sudden, with the average number of language nodes in a population increasing from one to five over the course of 30 generations or less.




Figure 5.5 A further 9 test runs with same parameters as Figure 5.3

This consistent sharp increase implies that this model supports the discontinuous view of language evolution (see discussion in Chapter 2). However, due to the genetic representation used, only one mutation is required to allow a jump from one language node to three or five nodes (Table 5.2). Once such a mutation, favoured due to the implementation details, has occurred the new gene succeeds quickly in the population. A second mutation might increase the number of nodes to seven, allowing for some improvement in language negotiation, and some of the results show this distinct stepping (Figure 5.5 top right, middle centre and, more clearly, bottom left). After the initial rise in the number of nodes, however, there is less significant selective advantage to be gained from further increases, and the graphs show some variation.

5.4 Modelling Language-Physiology Coevolution, Part II

The model has demonstrated the coevolution of the agent physiology with the agent communication schemes. To prove that the observed discontinuity is indeed a consequence of the model's implementation, the experiment must be repeated. Here integers are used directly in the genetic representation, rather than binary numbers.

As there is only one hereditary trait, the gene string for each agent consists of only a single integer. Different implementations are required for the crossover and mutation operators than those already presented, and these new operators are described next.

5.4.1 Crossover and Mutation

Inheritance works quite differently in the integer evolution strategy used from that in the previous model. Given two parents A and B, and two children C and D, each child will inherit its single integer gene string from a different parent.

$$\begin{aligned} \mathbf{C} &= \mathbf{A} \text{ and } \mathbf{D} = \mathbf{B} \\ \text{or} \\ \mathbf{C} &= \mathbf{B} \text{ and } \mathbf{D} = \mathbf{A} \end{aligned} \tag{5.8}$$

There is a 50% chance of either outcome. Thus, there is no crossover of genetic material at all.

As before, for each integer in a gene string, the mutation rate determines the likelihood of a mutation occurring. When a mutation occurs, the integer value is modified in some way, either incremented or decremented by 1, with an equal chance of either occurring. Unlike the outcome of inverting a bit in a binary representation, this form of mutation is potentially open-ended, with very large or even negative values possible. The results of mutation can be limited by the operator, which here forces any negative result to 0, and any value greater than 7 back to 7.

5.4.2 Parameters and Settings

Table 5.6 lists the parameters and settings for the second coevolution test.

Parameters for coevolution (using integer gene strings)			
Learning Rate, α	0.2	Mutation rate	0.015
Population	100	Training rounds, t	200
M	8	$d\alpha (\alpha / t)$	0.001
Standard Deviation	1.0		

Table 5.6 Parameters and settings

Again, the initial population of agents all have $N = 1$. The mutation rate is three times the previous value, but as the gene strings are now one-third the previous length, the number of expected mutations per generation has been kept constant. The nature of the mutations has changed, as described above.

5.4.3 Continuous Evolution of Signalling Ability

A typical run using the parameters shown in Table 5.6 produced the results shown in Figure 5.6. In this figure, the gradual evolution of language ability in the population is apparent. In this test run, over the course of 500 generations, the average number of active signal production neurons rises from 1 to 6, and the average success at interpreting signals from around 20 to over 80 percent.

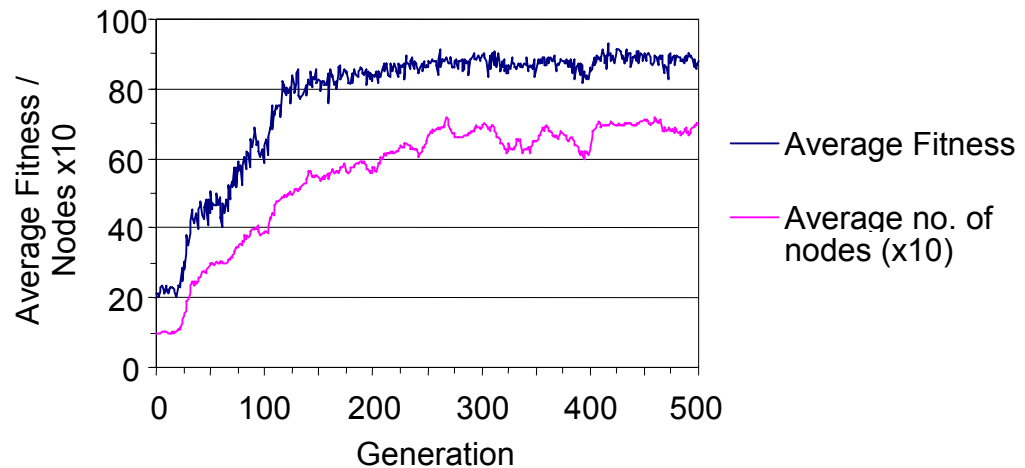


Figure 5.6 Coevolution of language and physiology. As before, as the communication capability increases (number of nodes), so does the success rate (fitness).

A further nine sets of results from runs of 250 generations using the same parameters and setup are shown in Figure 5.7. Despite individual differences the overall result of significant, yet gradual, improvement in communicative ability is consistent. Where large physiological changes are not possible due to single major mutations, it seems that a succession of minor adaptations spread through the population.

The coevolution of language and physiology has now been demonstrated using two different versions of the model. The models can be used to support either the discontinuous or the continuous EoL, and firm conclusions are hard to draw from these results.

Other settings and parameters can be varied however, to try to give some idea of the effects different natural conditions may have had on the evolution of language in nature. In particular, we can determine if spatial selection (our chosen approximation to kin selection) has a significant effect, by varying the neighbourhood size and repeating the experiment. This is done for the following test.

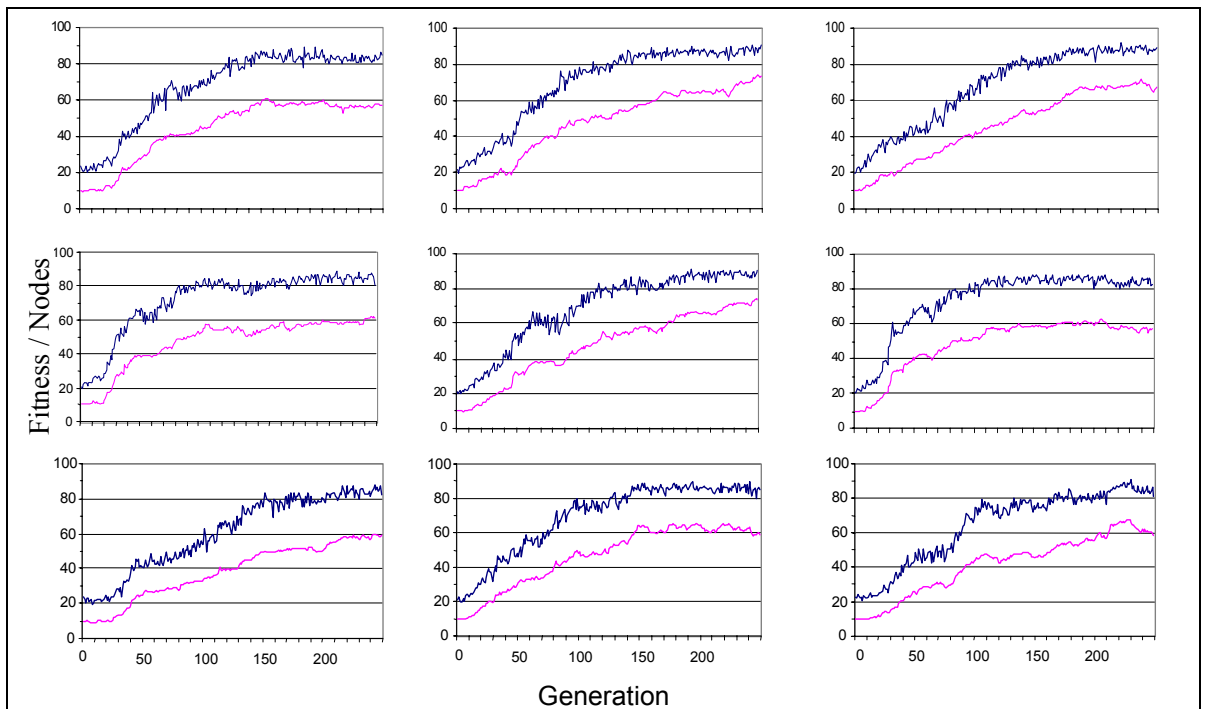


Figure 5.7 A further 9 test runs with same parameters as in Figure 5.6

5.5 The Effect of Neighbourhood Size

Spatial selection is one mechanism by which the likelihood of the evolution of cooperation may be improved. Like kin selection, such a mechanism can be responsible for the emergence of cooperative behaviour in situations where individuals may gain more by defecting than by cooperating.

In the model we have described, we can consider those agents which provide good and distinct signals (requiring a larger number of language nodes) to be cooperators, and those which provide poor signals (requiring fewer or even no language nodes) to be defectors.

If there is no benefit in defecting but some – even indirect – benefit of cooperation we should expect cooperation, as evidenced by success in communication, to succeed, even in the absence of spatial selection. Currently, defection does not benefit the agents in any way and so, even with a weakened spatial selection, cooperation should succeed in the population.

5.5.1 Parameters and Settings

This experiment is a repeat of the previous one (with integer representation), with only one change. The parameters are as listed in Table 5.6, except for the standard deviation which is set to 6.0.

5.5.2 The Effect of Large Neighbourhoods

We display the results of nine runs with large neighbourhoods in Figure 5.8.

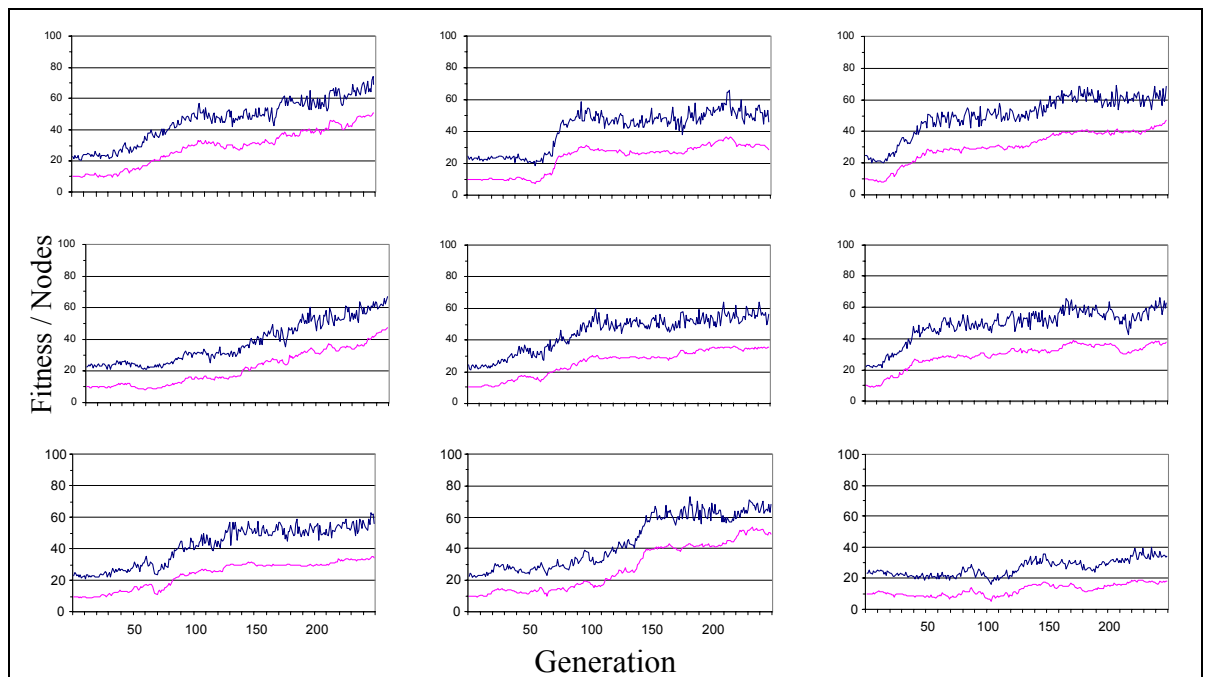


Figure 5.8: Coevolution with weakened spatial selection strength.

There is greater variation in these results than previously observed. What appears to happen is that the evolution of language is not prevented, but it is slowed down. This will be due to a decrease in the selective advantage of being a good signaller. As the model only rewards agents directly for successfully interpreting signals, the benefit a good signaller might receive is indirect. Such a signaller might improve the fitness of a neighbour, and in turn be selected as a mate for that neighbour if the neighbour has been selected. That neighbour, being a successful receiver, will have a greater chance of selection.

With larger neighbourhoods, an agent will receive few signals from a lone good-signaller within its neighbourhood. The signaller also has a lower chance of being selected as a mate by any of the agents it has provided with good signals.

Thus, the advantages of cooperation are lessened, and the evolution of language inhibited. But with no advantage in defection, it is seen that in most cases the communication abilities and success rates of the populations do eventually improve and progress towards the values seen in the previous two experiments.

The next step is to see the effect spatial selection has where there *is* a cost associated with being a good signaller.

5.6 Coevolution with a Costly Language Ability

As reviewed in Chapter two, there *are* costs associated with the ability to use language. Such costs can be added to the model by applying a fitness penalty to each agent dependant on the number of active signal production neurons, N , it possesses. Other than this simple change, no other modifications are required.

For the experiments performed here, the fitness penalty, $f(N)$, is:

$$f(N) = N^2 \quad (5.9)$$

This penalty has been arbitrarily chosen to penalise agents with larger numbers of active signal production neurons, to reflect the supposed adaptive costs of larger brain size. After evaluating an agent's fitness (Figure 4.6), the fitness penalty is applied. The resultant fitness score is used when selecting parent agents to form the succeeding generation.

5.6.1 Parameters and Settings

This experiment repeats the experiment of Section 5.4, with the addition of a fitness penalty as described above. An integer representation is again used in the gene strings. To test the effect of spatial selection, the experiment is repeated for the following different values of neighbourhood standard deviations: 0.2, 1, 3, 6 and 12.

5.6.2 Spatial Selection and Costly Language Ability

Results from using a standard deviation of 1 or 0.2 are shown in Figures 5.9 and 5.10 respectively.

These results are clearly poorer than shown in Figures 5.6 and 5.7. The fitness scores shown are plotted after the fitness penalty has been applied, and are consequently lower than the actual communication success rate achieved.

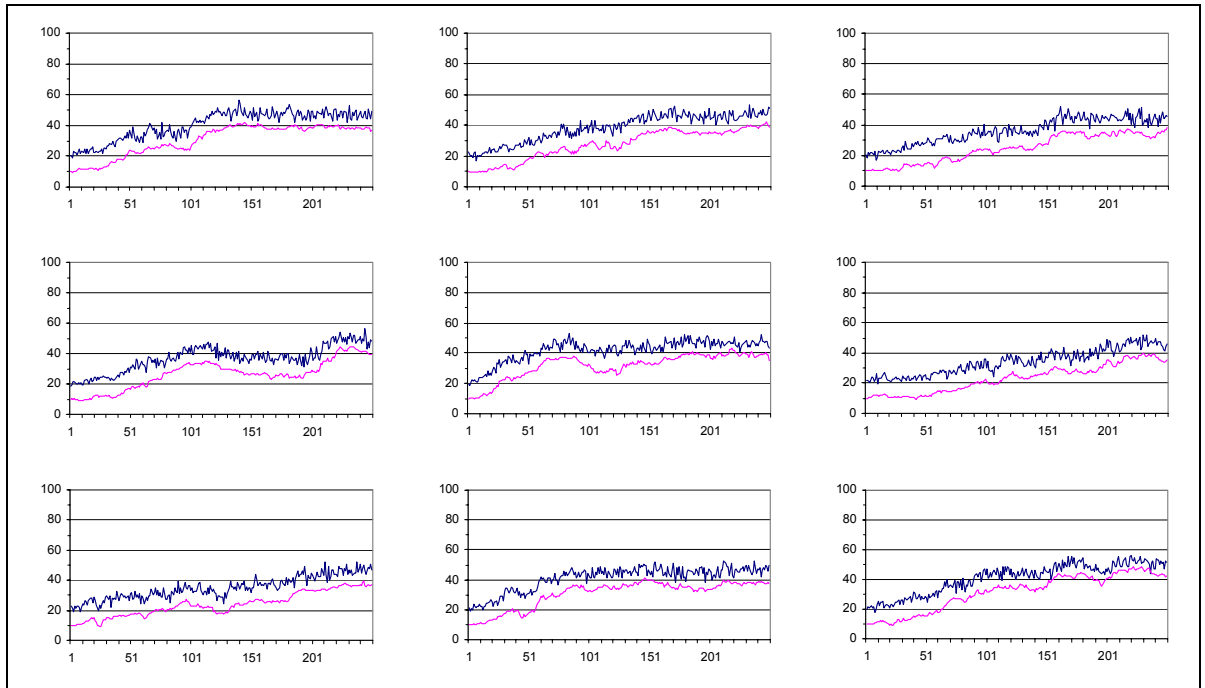


Figure 5.9: Coevolution with a costly language ability, showing average fitness and nodes $\times 10$ over 250 generations. Neighbourhood defined by a normal distribution of standard deviation 1.

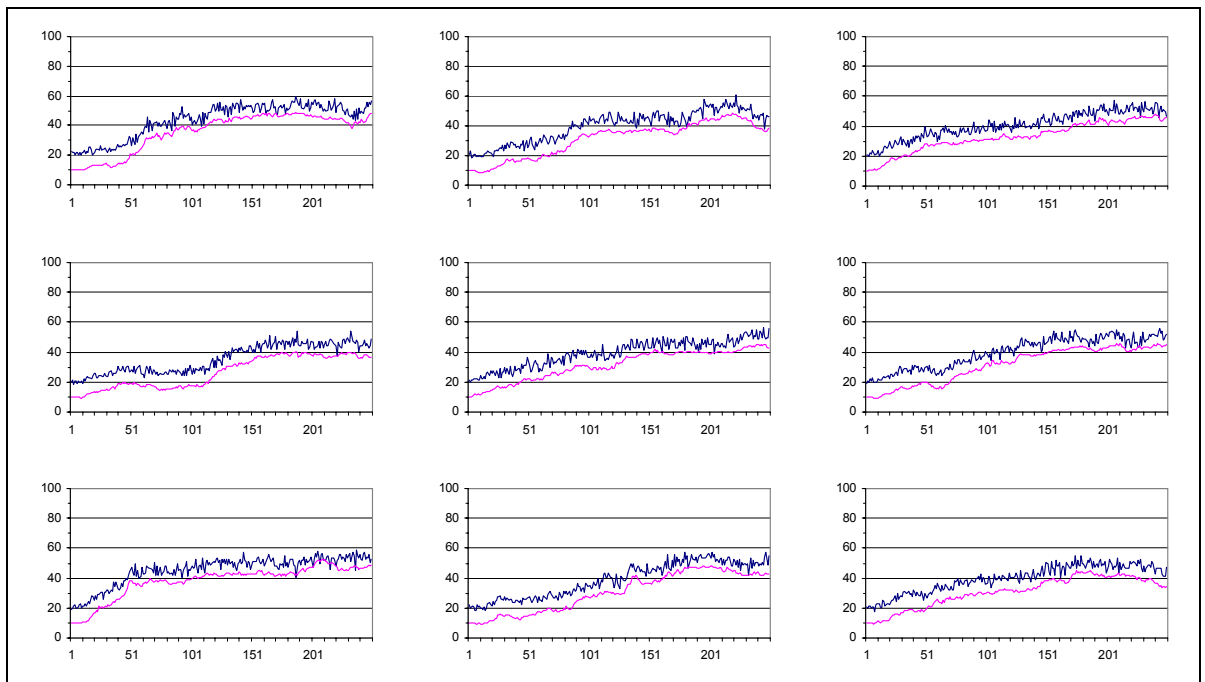


Figure 5.10: Coevolution with a costly language ability. Neighbourhood defined by a normal distribution of standard deviation 0.2. Showing average fitness and nodes $\times 10$ over 250 generations.

A slow and fairly steady improvement is observed in all eighteen sets of results, and there is little difference between the results for the two different values of standard deviation. The average success rate (not fitness) over all nine runs for each value of standard

deviation is shown in Figure 5.11. The minor difference between the results is shown more clearly. With the smaller neighbourhood size, the force of spatial selection is stronger and the evolution of linguistic ability enhanced slightly. This result may not be particularly significant however, as it is taken from just nine runs under each set of conditions, and may be influenced unduly by stochastic effects.




Figure 5.11: Average success rates for costly communication with neighbourhoods of standard deviation 0.2 (top line) and 1.0 (bottom line).

This seems to show that spatial selection does have an effect. For further evidence, we then repeat the experiment with increasing values of standard deviation. Figures 5.12, 5.13 and 5.14 show the results for standard deviations 3, 6 and 12 respectively.

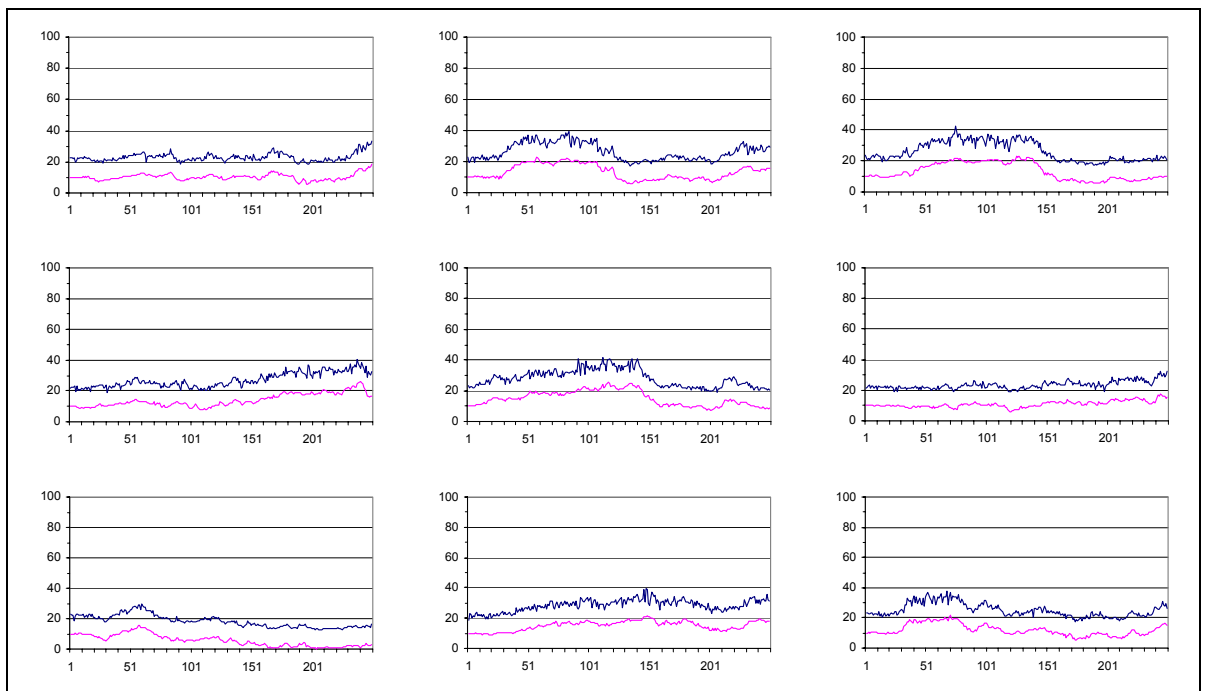


Figure 5.12: Coevolution with a costly language ability. Neighbourhood defined by a normal distribution of standard deviation 3. Showing average fitness and nodes $\times 10$ over 250 generations.

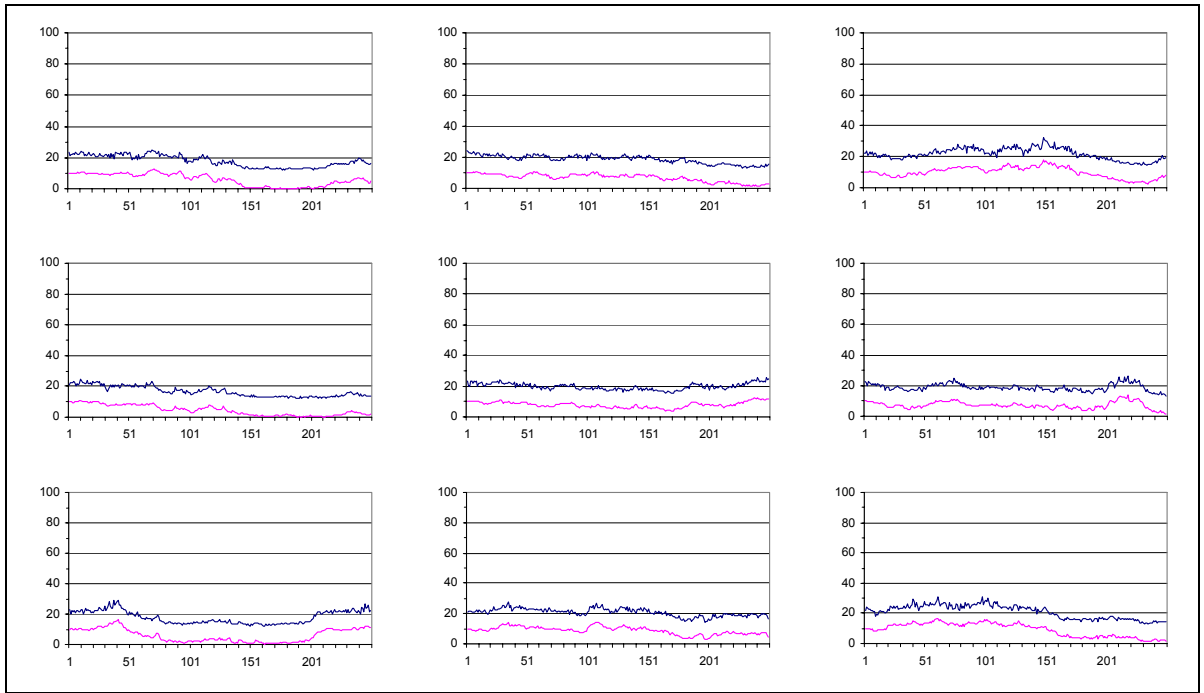


Figure 5.13: Coevolution with a costly language ability. Neighbourhood defined by a normal distribution of standard deviation 6. Showing average fitness and nodes $\times 10$ over 250 generations.

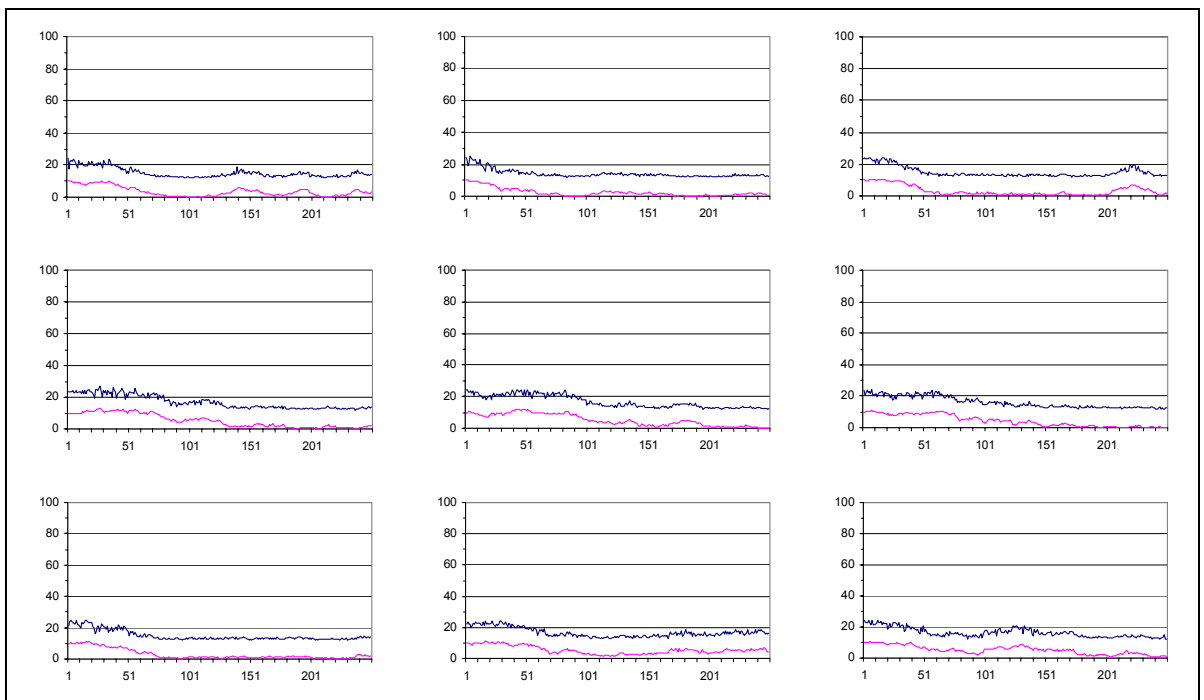


Figure 5.14: Coevolution with a costly language ability. Neighbourhood defined by a normal distribution of standard deviation 12. Showing average fitness and nodes $\times 10$ over 250 generations.

Reviewing the results presented in Figures 5.9 to 5.14, it can be seen that as the neighbourhood size increases the average linguistic ability that evolves get progressively poorer.

With small neighbourhoods a significant improvement is observed in the number of active language neurons and in the average communicative success rate. With very large neighbourhoods, Figure 5.14, the population is unable to even maintain the initial level of linguistic ability. This is despite a very small fitness penalty for having only one active language production neuron – only one point of fitness. Under these conditions the average communicative success approaches a pure chance level of success of 1 in 8 or 12.5%.

These results show clearly the strong influence that spatial selection could have on the evolution of linguistic ability. The cost of such an ability is borne by the individual speakers, while the benefits only arise from the interactions of multiple speakers. Spatial selection affects the model in a number of ways.

- During language negotiation, agents within the neighbourhood of good signallers will receive better signals, and can learn to interpret them successfully.

It is also not required that all the agents in the population use the same signals, and it appears that during negotiation local clusters of signal dialects emerge. This is investigated more thoroughly in the following chapter.

- During fitness evaluation, signals are received from agents within the neighbourhood of the agent under test. Thus the signalling ability or signals used by agents distant from the agent being evaluated do not influence the evaluation.
- During mate selection and mating, once an agent has been selected for reproduction by fitness, its partner will be selected randomly from those within the same neighbourhood.

The result of these effects is that the benefits a good signaller provides to its neighbours are more likely to be reciprocated – in that the neighbours will be more likely to use a compatible communication scheme, and will be more likely to pick the good signaller as a mate – where neighbourhoods are smaller. Thus the spatial selection can help overcome the fitness penalties suffered by altruistic individuals.

The simulation runs described in this section all used an integer representation for the agent genes (Section 5.4). Using a binary representation (Section 5.3) does not alter the outcome in any significant way (other than appearance of discontinuities, discussed further below). While not repeated here, the results of simulation runs using the binary representation are discussed in (Livingstone and Fyfe, 2000)

5.7 Discussion

5.7.1 Continuous versus Discontinuous Evolution of Language

The model in Section 5.3 demonstrates the discontinuous evolution of linguistic ability. Due to the implementation details of the model, the evolution of linguistic ability proceeds more by large jumps than by incremental steps. Making a small change to the model, the evolution of linguistic ability proceeds in small incremental steps, Section 5.4.

As noted in Section 2.4, there has been recent debate over whether the EoL has been the result of a single ‘macro-mutation’ or the result of a continuous process of gradual change and adaptation. The experiments presented seem to support each of these positions in turn. Rather than provide evidence to support one side of the continuity-discontinuity debate, the model appears to give some evidence against arguments that either is not possible.

However, the continuity-discontinuity debate appears to be settling down and largely resolved in favour of a position that accepts a large degree of continuity (Aitchison, 1998). Even some of the most notable opponents of the continuous EoL have in recent years modified their arguments considerably, accepting elements of continuous evolutionary theory into their thoughts (compare (Bickerton, 1984) and (Calvin and Bickerton, 2000)). The settling of this debate has also occurred with little regard to input from computational modelling based research. Our model demonstrates the successful use of communication in populations of heterogeneous language ability, but the scientific value of this alone is questionable. Whether the EoL was continuous or not, a time where individuals possessed different degrees of linguistic ability may have existed.

5.7.2 Spatial Selection

The model is perhaps more successful at demonstrating the need for additional selection mechanisms beyond natural selection to account for the EoL. That natural selection alone is insufficient in many cases to account for the evolution of altruistic behaviour is well known. Many studies have used theoretical or computational models to argue for the effectiveness of additional selection mechanisms in enabling the evolution of cooperation (several have already been cited, see Chapter 2).

With the assumptions that language can be costly to the individuals equipped with the ability to produce language, as previously argued, it is perhaps to be expected that additional selection mechanisms are required. We have demonstrated the EoL with spatial selection. It is also possible that other mechanisms influenced the natural EoL, such as kin

selection (Hamilton, 1964). See Di Paolo (1999a) for a detailed exploration of the application of kin and spatial selection in artificial life models.

Despite the existence of other models that show how spatial or kin selection can lead to the evolution of cooperation (as mentioned in chapter 2), we believe that this is the first demonstration of the effect of spatial selection on innate traits which have an acquired expression – an innate ability which can lead to cooperation only as the result of a learned behaviour.

5.7.3 Redundancy and Linguistic Ability

A further observation, which can be made when reviewing the results of the various experiments presented in this chapter, is that, in many cases, the populations evolve a redundant language capability. That is, the average number of language production neurons exceeds the number required to produce a different distinct signal for each of the possible meanings.

The minimum required number of language nodes for successful communication is just three. In contrast, in those experiments where language negotiation/evolution succeeds, the average number of nodes exceeds this – as shown in Table 5.7.

Experiment	Average number of nodes
Without cost, binary representation	6.2
Without cost II, integer representation	6.3
With cost, sdev = 1	3.8
With cost, sdev = 0.2	4.2

Table 5.7 The average number of nodes (over all runs shown) in the language layer in agents from the final generation of experiments shown in sections 5.3, 5.4 and 5.6.

That redundancy has an important role to play in language is recognised in linguistics. Without redundancy, signals would be very susceptible to disruption from noise in the environment, and too easily distorted (e.g. Crystal, 1987 p.146). This appears to be the main role of redundancy in existing explanations. Pinker (1994, p181) refers to Quine (1987) on this aspect of redundancy, who also argues that it provides a fallback or failsafe against disruption.

There is no noise present in our model however, and for redundant signalling capability to emerge repeatedly, even when there is a cost applied to the signalling ability, it would appear to be serving some additional purpose. There are in fact two such possible benefits of redundancy in the signalling capabilities of the agents.

First, redundancy may allow the agents to successfully interpret different signals which are used to represent the same meaning by different neighbouring agents. Lass (1997, Ch. 6) discusses this use of redundancy in language. ‘Linguistic junk’, as Lass terms it, allows speakers of differing dialects (or different variants of a common dialect) to communicate more successfully than would otherwise be possible as it may provide supplementary clues as to what is being said. We look at dialects in more detail in the following chapter, where we explore the effect that introducing neighbourhood limited communication has on the model.

The second, related, explanation is that redundancy makes the signalling scheme/language learning task easier and more likely to lead to success (evidence of this is found in the results in Chapter 4 where it can be seen that the addition of additional – redundant – language nodes improves the agents’ success at learning a common signalling scheme). That redundancy in representational capability can improve learning is known to researchers in ANN (e.g. Barlow, 2001), but does not appear to be widely recognised as an important benefit of redundancy in human language learning – and is not mentioned by any of the works cited in this section. This perhaps warrants further, future, investigation.

5.7.4 Investigating The Adaptive Benefits of Language

Throughout the experiments presented in this chapter we have made the, not insignificant, assumption that language is used cooperatively, ignoring the selfish applications of language. In doing so we have also abstracted the actual use of language considerably away from any specific application, simply assuming that it is beneficial for agents who are able to comprehend the signals produced by others. To justify this assumption necessitates a further review of some relevant literature. This concentrates on other ALife work, in attempting to demonstrate the particular problems that may occur when attempting to relax this assumption in agent based simulations.

First, the assumption is explicitly included by MacLennan (1991) in his innovative work on modelling the evolution of communication, where he limits his definition of communication to signals produced by individuals where the signal is of benefit (or, presumably, intended benefit) either to the signaller or to others of its group. While this is not part of the definition of language used in this thesis, the models here use a similar assumption, and reward listeners for correctly interpreting the signals they receive.

This simplification may appear questionable, but relaxing the assumption of cooperative communication is problematic, not just for computer modelling but also in more

theoretical approaches to the evolution of communication (Bullock, 2000). Including the possibility that communication may be used for both deceptive and cooperative purposes in a model of the evolution of language will lead to a complicated model, difficult to develop and to interpret. Some means of determining the honesty of different agents and the degree by which agents trust other agents would be required. This is far from the minimal models that we seek to build, and beyond the scope of this thesis.

In Chapter 2 we reviewed the work of Krebs and Dawkins (1984) that illustrates some situations that challenge the assumption the communication is necessarily cooperative in nature – such as manipulative signals that benefit the signaller but not the receiver. A possible resolution of this is to assume a distinction between cooperating in the development of a shared language versus actual cooperative behaviour (Grice, 1975; Lee, 2000). To be able to communicate at all, to cooperate or to compete, requires the cooperation in the distinct process of language negotiation. To some extent, this avoids rather than solves the problem that real or simulated agents should be capable of using language for non-cooperative ends. Whether agents are cooperating or competing, their ability to understand linguistic communication is indicative of their ability to use language for their own benefit.

5.7.5 Embodied Communication

Alternatively, we might accept that the assumption of cooperative behaviour is required to make the construction of our models tractable, but might not wish to reward agents for simply being able to understand the signals generated by each other. It may seem more realistic, and somehow better, to have models in which the real world benefits of language use are represented more explicitly than in the model which has been presented in this chapter.

Rewarding the successful replication of internal states – an abstraction of the ability to understand what message has been sent – has to be replaced by somehow rewarding the external behaviour of agents, this behaviour being affected by attempted communication. This requires ‘embodying’ the agents in some form of artificial environment to some extent (some examples follow). Success at tasks within the artificial world – possibly finding mates or gathering resources – leads to improved reproductive chances. Where the use of communication can lead to more successful behaviours, it should be possible for the communication strategies or abilities to evolve.

This approach has its own limitations as the following examples help illustrate.

Werner and Dyer (1991) evolved communication in an artificial ecology in which a communication protocol allowed immobile females to emit signals to guide blind males to them for mating. Mating and reproduction occurs when the task is successfully solved, rather than relying on external measures of fitness. The authors claim that the XGA – a genetic algorithm with mating the result of success in the artificial environment – is more realistic than arbitrary fitness measures. The lack of realism in the design of the model might challenge this idea, however. Additionally, signalling can serve only one purpose in this model – and this is predetermined, with tightly constrained sets of possible signals and responses.

The environment is such that matching signal/response pairs will lead to greater success at the mating task. There is no room for alternative strategies to evolve which might lead to greater success other than to communicate. Poor communication will not necessarily lead to failure however as a lucky male might still find a female anyway. But overall, this is the same as what is approximated by an external fitness function with roulette wheel selection, to the extent that it is not clear that any significant difference would occur by replacing the XGA with such a selection scheme. In both, successful communication leads to higher chance of reproduction, and in both it is possible for poor communicators to succeed nevertheless.

A work which combines the embodied evolution of communication with a non-embodied selection mechanism – in this case one based on elitism – is presented by Cangelosi and Parisi (1998), who evolves language in a population of ANN. The ANN based agents move round an environment in which there are mushrooms of poisonous and nutritious varieties. The task selected for language learning is for agents to be able to inform each other whether mushrooms are edible or are poisonous (as the ANN agents consume mushrooms they gain fitness rewards or penalties accordingly). Once all have had a period of testing in the artificial world some will be selected for reproduction according to their fitness. This model successfully demonstrates the evolution of signalling in an embodied model. Yet it is unclear what effect on the results, if any, is achieved by having the agents move around an environment. The learning task, and the decision to eat or reject mushrooms could equally well be undertaken without this extra level of detail, leaving an equivalent minimal model.

Despite the assertion of Werner and Dyer that embodied models are required for modelling the evolution of communication, there is a significant limitation, one that is present in these models. In these the communication is tied to a particular behaviour,

chosen by the experimenter. If we are attempting to explain the evolution of language, then this could be as serious a problem as *not* having any specific adaptive benefit from the use of signals. Demonstrating how language might have evolved to serve a particular function borders on presenting a ‘just-so story’, proclaiming the original reason for the evolution of language, even though this may not be the intent of the author.

Further, the constraints in such models are such that by providing only one possible benefit of successful communication, and by enabling the emergence of language, the result is largely pre-determined and predictable: that language/communication can evolve to provide the stated benefit. Negative results provide a useful contrast and, as noted by Hurford (1992), are required to highlight what the requirements are for positive results, but such failures to evolve communication or language are missing from many of these models.

One reason for the lack of negative results is that there is no cost associated with communication or cooperation. If communicating is free, and cooperation brings no loss of fitness but successful cooperation brings benefits then the evolution of communication or language is to be expected – particularly given the constraints placed on the range of possible evolutionary changes.

Accepting the great complexity of arguments over the origin of language ability in humans (as reviewed in Chapter 2) and the very many functions it serves, embodying language such that it provides a single particular benefit is no better than making the sweeping assumption that language is of benefit to those who use it (to transmit or to receive information). Such embodied models provide a host of alternative possible explanations for the evolution of language with little way of choosing between them.

The concept of building a model which allows agents to evolve communication to fulfil a wide variety of uses, which individually and combined provide adaptive benefits but are not predetermined, is not inconceivable. It is, however, considerably more complex than any model which has been implemented to date. Attempts to embody communication in an environment without a predefined role or purpose have to date had little success (e.g. Werner and Dyer, 1993).

This relates to problems in using ALife as a means to understanding the evolution of self-replicating systems as noted, at some length, by Taylor (1998). The environment and the rules governing interaction and replication provided by the computer program constrains the possible evolution that can occur, even where the results surprise the experimenter

who built the model. To evolve language in ALife, a model must first be built in which language can evolve.

The lack of clear benefits of embedding models compares poorly with the advantages of not embodying the communication, which can be significant in time spent developing and running simulations. As argued, a lot of what occurs in some embodied models is simply ‘window dressing’, which, stripped away, may leave an equivalent, simpler, model without embodiment or the attached overheads of such modelling.

So, while it may be attractive to remove as many assumptions as possible, and to embody a simulation of evolution within a ‘realistic’ model that includes an environment modelled on reality, doing so does not necessarily improve a model. With the additional complications that arise it is quite reasonable to fall back on some common assumptions. In modelling the evolution of language it is possible to reward successful use of language without holding that all language is cooperative or believing that merely understanding language confers a benefit. But such simplifications can be built into models without necessarily compromising the models.

5.7.6 Limitations and Shortcomings

As should be clear, the successful application of computer models is far from being a panacea for the difficult problems surrounding the understanding of the EoL. To be best able to use such models constructively it is important to appreciate their limitations.

As noted above, computer models are good for demonstrations of how different uses of language provide benefit – but they are not much good for proving why and how language evolved. In the EoL in *Homo sapiens*, historical accident may have played a large role, and many distinct adaptations to different environmental pressures appear to have had some role in pre-disposing our hominid ancestors to language (Deacon, 1997).

Over the course of time a large part has been played by historical accident and serendipity – what makes the evolution of language interesting to scholars is that it occurred in the evolution of *Homo sapiens* alone amongst all the species that exist in the world. With such a large part played by chance, it is unlikely that a computer model will ever conclusively demonstrate the precise conditions that led to the emergence of language. A better understanding is more likely to come from the rare finds of fossils of our hominid ancestors, from which much of the current knowledge about the evolution of language, and the required physiology to support it, is derived.

A particular shortcoming with the model used in this chapter for modelling the EoL is that only very limited evolution is possible, with a small and closed set of possibilities. A consequence of this is that it is easy for language to emerge in this model when the costs favour its emergence, but it is much harder in nature where the search space is open ended.

5.8 Summary

In this chapter the experimental model described in Chapter 4 was used in a series of experiments investigating the co-evolution of language and physiology. The first two experiments show populations of agents evolving to be capable of using better signalling schemes. As argued, this is the first known demonstration of the evolution of communication in populations where the capability to use learned signals evolves, rather than simply the signals themselves, or sets of initial weights. Then, comparing the models it is clear that depending on the details of implementation, models may be constructed which support arguments on both sides of the (now somewhat settled) continuity-discontinuity debate. This highlights some of the problems discussed in the previous chapter, and expanded in this one – of how the implementation of an ALife model can impact on the results in undesired ways, and of how different implementations of the same model might give quite different results. In this case, different ways of representing genes themselves leads to contradictory outcomes.

More satisfying is the subsequent demonstration of the requirement for some additional means of selective force – such as kin or spatial selection – for the successful emergence of the ability to use language. It is known that language both incurs many additional costs as well as serves many functions – and this work supports theories that see the social nature of hominids as being key to subsequent evolution of language. Where a number of authors point to the social functions that language serves as being responsible for its emergence (Chapter 2), this work dovetails with the complementary notion that without suitable social groups existing the fitness costs could easily prevent the emergence of language.

Another observation was that where language evolution was successful, the agents evolved such that they had redundant language capabilities. While some authors have noted the different roles of redundancy in language, it is generally ignored in most work on the EoL. The spontaneous emergence of redundant linguistic capabilities in our model, and the effect it has on agent fitness, emphasises the importance of redundancy. This significant observation warrants further, future, investigation.

Finally, it was also noted that different dialects of signalling schemes emerged in populations where neighbourhoods limited the communication interactions. This spontaneous emergence of dialects is investigated more thoroughly in the next chapter where attention is turned to the cultural evolution of languages.

Chapter 6 The Cultural Evolution of Language and The Emergence and Maintenance of Linguistic Diversity

6.1 Introduction

Despite the somewhat negative conclusions from the previous chapter, there is one interesting result that can be observed. A listing of the signals produced by the different agents in the population shows that not all agents are using the same signals, as shown in Table 6.1, which displays some of the signals generated by a sub-section of the population from one experimental run. The full table, showing all signals produced by all agents is included in Appendix A.

Agent	Nodes	Meaning 1	Meaning 2	Meaning 3	Meaning 4
0	6	+ + + + - +	+ - - - - +	- + + + + +	- - + + - -
1	6	- + + - - +	+ - - - - -	- + + + + +	- - + + - -
2	6	- + + - - +	+ - - - - -	- + + + + +	- - + + - -
3	6	- - + - - +	+ - - - - +	- - + - + +	- - + + - -
4	5	- - + - -	+ - - + -	- - + - +	- - + - -
50	6	- - - - - +	- - + + + +	+ + + + - +	- - + - - +
51	7	- + - + - + +	- - + + + - -	+ + + + - + +	- - + - + + +
52	7	- - - + - + +	- + + + + - -	+ + + + + + -	- + + - - - +
53	7	- - - + + - +	- + + - + - -	+ + + + + + -	+ - + - + - +
54	7	+ + - + - + +	- + + - - - -	- + + + + + -	+ - + - - + +

Table 6.1: Reviewing the signals from one set of results from one of experiments shown in section 5.4

While adjacent agents use very similar signals, there appear to be greater differences between the signals used by non-neighbouring agents. It appears as if signal ‘dialects’ have emerged (also see Section 4.9.1). This is contrary to the previously reported results of Nettle (1999a) and arguments of Milroy (1993) (see section 2.5.2), which hold that only where diversity provides an adaptive benefit will it emerge. There is no benefit for using a different signal than a neighbour’s in this model, however. On the contrary, a situation where all agents use the exact same signals would intuitively provide the maximal benefit to the agents.

A reformulation of the previous model to allow the evolution of signal diversity to be studied in isolation, without any possible influence that the evolution of the network structure may exert, is required to examine this more closely.

This is presented in this chapter, and we find that it supports the conclusions drawn from the initial observation, prompting a number of experiments to examine this phenomenon.

This provides some evidence for the possibility that much of human linguistic evolution is adaptively neutral, and arguments regarding this are also presented and reviewed in more detail.

6.2 An Initial Test

The first reformulation forms a single homogeneous population of agents, and allows them to negotiate a communication scheme over a number of training rounds. Only one generation is used, and the training algorithm for language learning is unaltered from the previous chapters.

A homogeneous language population with 4 language units is trained for 5,000 rounds, starting from random signals. The long training time allows a very high degree of coordination amongst the agents but, due to the local communication, does not necessitate a common language over the whole population. It is observed that large neighbourhoods negotiate a common signal for a given meaning, but distant agents may have significant differences in communication schemes used. At the boundaries between neighbourhoods, agents may exist which interpret signals from different schemes correctly. This is shown in Table 6.2. The three agents included each attained a maximal fitness score, interpreting all signals correctly.

Meaning 0							Meaning 4						
Agent	11	1	-1	1	-1	1	Agent	11	1	1	-1	-1	-1
Agent	12	1	-1	1	-1	1	Agent	12	1	1	1	-1	-1
Agent	13	1	-1	-1	-1	1	Agent	13	1	1	1	-1	-1
Meaning 1							Meaning 5						
Agent	11	1	1	-1	1	-1	Agent	11	1	1	-1	-1	1
Agent	12	1	1	-1	1	1	Agent	12	1	1	1	-1	1
Agent	13	1	1	-1	1	1	Agent	13	1	1	1	-1	1
Meaning 2							Meaning 6						
Agent	11	1	-1	1	1	-1	Agent	11	1	-1	1	1	1
Agent	12	1	-1	1	1	-1	Agent	12	1	-1	-1	1	1
Agent	13	1	-1	1	1	-1	Agent	13	1	-1	-1	1	1
Meaning 3													
Agent	11	1	-1	-1	-1	-1							
Agent	12	1	-1	-1	-1	-1							
Agent	13	1	-1	-1	1	-1							

Table 6.2: The signals (including bias) used by three adjacent agents for seven environmental states. All three scored maximum fitness, interpreting all signals correctly, despite differing communication schemes.

This appears to show the emergence of different signal dialects within the population. If this can be compared to human dialects, then it is possible that the model can be used to make inferences about the causes and formation of human dialect diversity. However there are no obvious reasons why dialects should form in this model, and there is no motivating factor to encourage the formation of dialects. This contradicts some of the views reviewed

in section 2.5 and directly challenges the results of Nettle (Nettle and Dunbar, 1997; Nettle, 1999a; Nettle, 1999b) that demonstrate the need for a social function for the successful emergence of dialect diversity.

It is possible to draw on these results to refute Nettle's claims and with further study, a stronger case can be made, but the model as applied in the previous chapter is not entirely suitable for several reasons. Firstly, it is difficult to visualise the results – large tables of data are required to review the dialects formed in one generation, and it would not be easily possible to compare the dialects over a number of generations.

Also, there may be some interference between the evolution of the agents' linguistic ability and languages learned – a few individuals with more limited language layers than their neighbours may effectively hamper communication and learning between more capable surrounding agents. Linguistic boundaries may be so formed, skewing the experimental results.

Therefore some changes are necessary to remove any possible unwanted effects of having structurally non-homogeneous populations and to improve the transparency of the model. The core details of the model remain the same, however.

6.3 Human Linguistic Diversity

6.3.1 Patterns of Diversity

Linguistic diversity has already been discussed to some extent in Chapter 2, but a further overview is called for before proceeding. Only the briefest of summaries is presented here; more extensive descriptions can be found in the literature (e.g. Crystal, 1987; Trudgill, 1995; or Chambers and Trudgill, 1998). Many of the issues relating to the EoL do not apply, and language diversity is itself a facet of human language that has been studied extensively, both with and without relation to the EoL.

Crystal (1987) presents an overview of some key characteristics of human language diversity. One important finding of dialectology is that the boundaries between different dialects or languages are not always easy to define. Geographically close dialects may have sufficiently similar grammars or lexicons to allow speakers from each dialect to understand each other quite well, despite differences in the dialects. Geographically distant dialects may not be at all mutually intelligible, despite being at either end of a chain of dialects where every dialect is intelligible to speakers of the neighbouring dialects, Figure 6.1. Several such dialect continua exist in Europe, blurring the boundaries

between the different languages within the Germanic, Scandinavian, Romance and Slavic language groups.




Figure 6.1: A schematic dialect continuum from dialect A to dialect E, showing some degree of mutual intelligibility between adjacent dialects (After Crystal, 1987, p25)




Figure 6.2: Dialect boundaries (After Crystal, 1987, p28)

6.3.2 Linguistic Boundaries

The geographical boundaries between dialects may not be easy to determine, and these boundaries themselves influence the future change and diversification of languages. Dialects may differ in their lexical, semantic, morphological or phonological features. Sampling the language use of individuals in some area results in a map, on which boundaries may be drawn to show where language use is distinct on either side according to a particular linguistic feature. These boundaries are termed *isoglosses*.

It may be expected that these lines will be largely coincident, forming clear dialect boundaries. Often, however, the boundaries are not even nearly coincident, as individuals near a boundary may use differing mixtures of lexical and grammatical items from the surrounding major dialects, Figure 6.2. Only when viewed at a more distant scale is it possible to determine distinct dialect areas – with a poorly demarcated boundary between them.

As if this picture were not complicated enough, it is also widely recognized that no two individuals use language in the exact same way – in a sense every individual speaks their own particular dialect, or *idiolect*. What is viewed as a dialect is merely some norm

derived by sampling many idiolects. This, as we shall see, is highly significant when examining or attempting to explain linguistic diversity.

Sharp linguistic boundaries that do exist often coincide with significant geographical boundaries, such as mountain regions, or with strong cultural boundaries (c.f. Chambers and Trudgill, 1980, p111). Where such boundaries exist, limiting the interactions of individuals across the divide, there may be many coincident isoglosses splitting the sides and a resultant lower degree of mutual intelligibility.

As well as linguistic and other cultural evidence that some cultural or physical barrier really does inhibit and limit interactions, genetic evidence can corroborate this. For example, studies in Africa have shown that certain genetic markers act as good predictors of the language group of a subject's spoken tongue (Renfrew, 1998). Where the boundaries limiting interaction are maintained by linguistic rather than cultural distinctiveness, language is itself the boundary working to maintain linguistic differences.

6.4 Analytical Models of Linguistic Diversity

While the features and characteristics of language diversity described in the previous section are well known, it is the nature of the mechanisms which give rise to them that we hope to illuminate and explore with the use of artificial life based models. Before we proceed, we will consider the contribution of analytical techniques using mathematical models, reviewing a few which have been developed for this purpose.

Like simulation models, mathematical models require some amount of abstraction and simplification (Chapter 3). A mathematical model of the evolution of language diversity in a large, spatially distributed, population with a large number of linguistic variables would not be tractable, and has not been attempted in any of the models which are reviewed here. The models that are possible are more constrained. For example, Pagel (2000) presents two models of interest. The first models the growth of linguistic diversity over time. The unit of this model is 'a language'. No notion or representation of linguistic features, population size or growth or geography exists in the model. While the model appears to describe the growth in the number of languages over time, no information about how or why languages evolve may be gained.

The second part explores the different number of languages that exist within closely related language groups. Maths is used to determine whether the difference is statistically significant. But for an explanation, Pagel resorts to means outside of mathematics. The analytical approach has here been used to *describe* language evolution, and a

demonstration made of how appropriate maths might be used to highlight a particular spread of language as being unusual and requiring further, non-mathematical, explanation.

6.4.1 The Niyogi-Berwick Model

Perhaps more ambitious is the model developed by Niyogi and Berwick (1995), henceforth the NB model, and generalised in Niyogi (2002) – which additionally relates the NB model to the Cavalli-Sforza and Feldman (1981) model of cultural evolution – to include spatial organisation of the language users. In order to understand the latter, a fairly detailed review of the former is first required. As well as being highly relevant, this model has been widely reviewed and accepted (e.g. Lightfoot, 1999, p102), justifying the extended treatment it is given here.

In Niyogi and Berwick (1995) a dynamical systems model of grammatical change is presented. In this model, two grammatical variants compete in a population. This model works on the basis that children are exposed to sample sentences produced by their parents, and that the children then acquire their own language grammars according to the sample presented and the learning algorithm used. It is then possible to derive mathematically the progress of competition between two grammars.

The actual competing grammars are taken from human languages, and may vary in different factors, such as head-first or head-last or verb position. For each variant, the set of possible grammatically correct sentence structures is generated. For example, the set of possible sentence types of the -V2 grammar (non-verb second structure of modern English, as used as an example in Niyogi (2002)) is shown below.

$$L_1 = \{ S V, S V O, S V O1 O2, S Aux V, S Aux V O, S Aux V O1 O2, Adv S V, Adv S V O, Adv S V O1 O2, Adv S Aux V, Adv S Aux V O, Adv S Aux V O1 O2 \}$$

Figure 6.3. The -V2 grammar corresponding to modern English as represented in the NB model, which uses only degree-0 sentences (without recursion or additional sub-clauses). S = subject, V = verb, O1 = direct object, O2 = indirect object, Aux = auxiliary, Adv = Adverb

For any two competing grammars, a number of the sentences may be common to both, while others will only be parsable in one or the other. The common sentences are considered ambiguous, in that speakers of either grammar may produce them. The Trigger Learning Algorithm (Gibson and Wexler, 1994), TLA, is the learning algorithm used. In this, individuals start with a randomly selected grammar, which is used unless a sentence which cannot be parsed in it is encountered. When this happens, an attempt is made to

parse the sentence with a different random grammar, and if this succeeds the new grammar is retained.

In a competition between two grammars L_1 and L_2 , if L_1 has a high proportion of ambiguous sentences, a , which are also valid L_2 sentences, while L_2 has few ambiguous sentences, b , then L_2 will be favoured, and over time the proportion of the population using L_2 may rise. This relies on the assumption that all sentence types have an equal likelihood of occurrence. As the ambiguous sentences are those that are valid in both grammars, the proportion of ambiguous sentences only differs where the grammars have different sized sets of possible sentence types – and the grammar with the larger set will have the smaller proportion of ambiguous sentences. In all cases, the competition evolves to a single fixed point, and this is derived for the situation where language learners learn from exactly two randomly selected examples, and in which there are an infinite number of examples.

If $a = b$, over time the competition will reach the fixed point of exactly half of the population using each variant. If $a < b$, the fixed point will be one in which L_1 is used by over half of the population or used by the *entire* population where an infinite number of learning examples are presented to learners. The NB model is successful in reproducing the logistic, s-shaped curve of language change, often observed in linguistics (see Aitchison, 1991).

Niyogi (2002) then generalises the model so as to include spatially distributed populations. It is this that is of particular interest here, where we are considering the evolution of dialect diversity. However, to make the basic maths tractable the spatial model assumes that all the speakers of L_1 are in one grouping adjacent to the L_2 grouping, and in every generation this separation is enforced. The outcome of the model is similar to that of the previous model, with the competition between the two dialects evolving to a stable fixed point. Due to the neighbourhood model used, languages are never completely eliminated however, as at the extremes there are always children who are not exposed to the other competing language. An alternative neighbourhood model, which does not arrange the speakers into the two homogeneous neighbouring groups, is postulated but the equations for this are not derived.

6.4.2 Criticism of the Niyogi-Berwick Model

Some weaknesses of this model are highlighted in Niyogi (2002) itself. In particular the unlikely possibility that the more complex and unusual sentences in L_1 occur with the

same frequency as the simple ones is central to the derivations and progress of competition between two grammars in this model. Attempts to apply the model to real instances of grammatical change seem problematical if this assumption proves to be inaccurate.

Further problems are described by Briscoe (2000c), who highlights some of the problems of using an analytical macro-model of a population rather than a stochastic, computer based micro-model. Using a stochastic simulation model, Briscoe argues that we should expect different results from the NB model – at least until the population becomes very large. In a micro-model stasis, such as represented by the fixed point end result in the NB model, is very improbable. Briscoe argues that, as yet, no clear advantage to macro modelling for *realistic* G (class of grammars), A (learning algorithms) and P (probability distributions) is demonstrated by the NB model. The key weakness of the NB model is in abstracting away the sampling issues – and is particularly bad for small groups. Additional important factors, such as movement, birthrate, proportions of speakers and the resultant linguistic mix of population, are either over-simplified or not represented at all in the model, and these can be represented more easily in stochastic than in analytical models.

Additionally, Clark (1996) shows that the logistic curve found in the NB work appears only where the selection is between exactly two grammars, a failing noted in Briscoe (2000c). Later in the chapter we will review the alternative stochastic models presented in Briscoe (2000a).

That the NB model always evolves to a fixed point, from which no further change or innovation will occur is obviously unrealistic. One of the claims in Niyogi and Berwick (1995) is that the NB model can be used to evaluate theories of grammatical acquisition (such as the TLA). It is unlikely that the NB model can satisfy this claim, requiring as many assumptions as it does, being limited to competition between two, and only two, distinct grammars, and lacking much of the richness of human interaction.

The generalised spatial model has some of its own problems. A poor grammar will never be eliminated completely – even though almost all of the population has converged on a different grammar. Such persistence is unlikely in general in real languages. Equally unlikely in real languages, a ‘better’ grammar with one speaker in a population of millions will eventually take over the population – leaving the previous common language a tiny minority tongue. This emphasises that although the model has been generalised to include a spatial dimension it completely fails to represent a *language ecology* (see Section 2.5.5) in any meaningful way.

6.4.3 The Cavalli-Sforza and Feldman Models

As noted in the previous section, Niyogi (2002) relates the NB model to the Cavalli-Sforza and Feldman (CS-F) model of 1981. Their analysis shows that both models have directly comparable outcomes. Despite minor variations, the results of both are qualitatively the same. No further analysis of the model is presented.

Here we present our own analysis using a simpler, earlier CS-F model, (Cavalli-Sforza and Feldman, 1978), to illustrate how the use of analytical models can lead to suspect conclusions. In this earlier model, the cultural traits acquired by an individual are determined by a weighted sum of the traits existing in the previous parent generation. The cultural trait, X , acquired by the i th agent in a population at time $t+1$ is determined by equation (2-1) (Chapter 2).

So the cultural traits acquired by any member of a population is determined by the cultural traits of those around them – weighted according to the amount of contact between individuals. Cavalli-Sforza and Feldman (1978) present some further analysis showing the qualitative effects of different transmission models and assumptions (one or two parents, and varying the amount of influence from other members of the parent generation) on the degree of variation that survives in the model. In a one parent model where each child learns only from their parent, it is noted that the variance will grow linearly over time and that eventually there would be no cultural homogeneity. That influence from other members of the population may be the mechanism which prevents this is postulated, but not shown analytically (other than for the case of two parents influencing the traits learned by children).

In applying this model to the eol we are interested in the case where the population around a learner, not just the learner's parents, assert influence over the acquisition of cultural traits – in this case language. We can say that the grammar acquired by the i th member of some population will be:

$$g_{i,t+1} = \sum_{j=1}^N w_{ij} g_{j,t} + \epsilon_i \quad (6-1)$$

$$\text{where } 0 \leq w_{ij} \leq 1 \text{ and } \sum_{j=1}^N w_{ij} = 1$$

This is analytically intractable if the weights w_{ij} vary for every pair of individuals within the population. If $0 \leq w_{ij} \leq 1$, then two agents i and j might exert a very strong influence on one another or may be entirely disconnected. Arbitrary values of w_{ij} could include

situations where a single population actually contains two distinct and unconnected populations. While this example is somewhat extreme, a random allocation of weights would more likely result in some highly connected clusters of agents, and other less well connected agents (similar to the social networks presented by Milroy (1980)). Simplification is required if these equations are to be solved mathematically. If we assume uniform contact amongst the whole population, despite the implausibility of such uniformity, we can lose the w term from the equations altogether. This simplifies the equation considerably:

$$g_{i,t+1} = \frac{1}{N} \sum_{j=1}^N g_{j,t} + \varepsilon_i \quad (6-2)$$

It should be clear that unless ε_i is large, the acquired grammars for any two agents could be only trivially different. A very strong averaging force is in effect, and as population replacement proceeds, variation from the average grammar would be small and would only survive at all as a result of the small errors, ε_i , introduced in learning.

6.4.4 Summary of Analytical Models

The analytical models discussed have some important differences, but possess underlying similarities in their treatment of the eol.

In attempting to provide a mathematical model of complex phenomena, many real world issues are simplified or removed altogether from the models. Spatial organisation is lost in most of the work, treating the populations as inhabiting the one common space.

The NB model is used to examine what happens when there are two competing grammars, and comes to some questionable conclusions (6.4.2). It also shows evolution favouring grammars with a grammatically selective advantage. This implies that selectively neutral changes cannot ultimately succeed in a population. The model also limits the possible eol as the population can only choose between two grammars – the grammars themselves are unable to change or evolve.

The presented CS-F model (6.4.3) presents an equation for how an individual trait possessed by a single individual within a population may be determined according to cultural influences – where the traits themselves exhibit no selective advantages. This allows each member of a population to possess a large number of different traits. In simplifying the model such that, for any initial distribution of trait values, it would be possible to calculate which particular values succeed in the population we find that this results in a model in which the population always converges on the average trait value.

This result showing that diversity would be expected to disappear from the population is not a necessary consequence of the model itself – but is due to the additional assumptions introduced to make the maths tractable. And indeed, the results of the work on the NB model may be the result not of the model itself, but consequences of the many additional assumptions and simplifications introduced to make the model solvable (Niyogi, 2002).

6.5 Modifying the Artificial Life Model for Exploring the Emergence of Linguistic Diversity

Having reviewed some of the attempts to mathematically model linguistic change and diversity, and highlighted some of the weaknesses of the approach – principally centred round the absence of any richness or randomness in the interactions of individuals – we proceed to extend the Artificial Life model presented in the last two chapters. Several modifications are required to allow an exploration of change and diversity, and these are explained next.

The most important change made is to remove the evolution of linguistic ability from the model. This simplifies the model, allowing more focus on the evolution of linguistic diversity without any interference from evolving or heterogeneous language abilities. These could otherwise distort or confuse the results by providing an additional potential source of linguistic diversity. Any linguistic diversity which emerges in the altered model can then not be due in any part to the influence of the EoL. This also results in a more realistic study of the effects of cultural evolution, as the eol works on a much faster time scale than the EoL.

Removing EoL from the model is achieved simply by fixing the number of active language production neurons, N , for all agents to a single arbitrary value. The selection of this value can help overcome another problem with using the model as previously described. It should be apparent, considering Figure 6.1, that comparison of the different signals used by a whole generation of agents will not be simple. Even less simple would be an attempt to compare signals used by agents over a great many generations. Some means of easy visualisation of results is required. Setting the value of N to 3 can enable visualisation.

Each signal will be exactly three bits long, allowing a repertoire of eight distinct signals in total. Using the individual bits, which compose a signal, to set red, green and blue colour values, each signal generated by an agent can be represented by a distinct colour.

Colour Plate 6.1 lists the colours that represent each of the eight possible signals. To obtain these colour values, the bipolar signal vectors are transformed to binary vectors – with all -1 values being reset to 0. For the remainder of this chapter, we treat the signal vectors as being binary vectors although this is solely for convenience for the consequent ease of mapping to colour values.

Allowing for redundancy in the signal \Leftrightarrow meaning mappings, the number of possible meanings, M , and signal bits, N , have to be set such that $M \ll 2^N$. As previously explained, for purposes of visualisation, $N = 3$, thus we require that $M \ll 8$. M was arbitrarily set to 3.

The model parameters and settings are summarised in Table 6.3.

Parameters for coevolution (using standard binary representation)			
Learning Rate, α	0.2	Training rounds, t	See below
Population	120	$d\alpha (\alpha / t)$	$0.2 / t$
M	3	Standard Deviation	1.5
N	3	Mutation Rate	0

Table 6.3 Parameters and settings

With no mutation, the population maintains a homogeneous signalling ability, but otherwise the implementation of individual agents is unaltered. With no evolution of the signalling ability, there is also no selection of individuals for reproduction. Success of communication can still be measured, but is no longer used as a fitness measure for mating. It is no longer the biological evolution that is of concern, but a cultural one. Within the model, child generations are created automatically, but signals are no longer learned simply from peers. Communication and learning across generations is required to investigate the eol. The cross-generational learning of signals is discussed in more detail in the next section.

6.6 Experiment 1: Emergence and Maintenance of Dialects

The first experiment is simply to verify the observation previously made – that dialects emerge in the model as a result of the learning interactions between the agents.

6.6.1 Experimental Setup

Within this and the following experiments, it is required that, during learning, training examples are provided by a parent generation to allow the signal repertoires to be learned over time. Agents are selected from the parent generation according to spatial distance from the learner. An agent occupying the same position as the learner, but in the parent

generation is at distance 0. Again, selection is based on a normal distribution centred at this point.

In human society language is learned from peers as well as elders, and this can be included in the model, by providing a set number of training examples from the parent generation, and a number from the current peer generation. While this makes the model more ‘human’ it does not necessarily follow that it will have any significant impact on the results.

This experiment was run with two model configurations: one with learning from the parent generation only, and one with learning from the immediate parent generation *and* the current (peer) generation. Table 6.4 details the number of training rounds applied in each case. All signals are initially random before training.

Additional Parameters			
Parent and peer learning		Learning from parents only	
Parent generation training, t_p	40	Training rounds, t	40
Current generation training, t_c	20		

Table 6.4 Parameters for parent and peer learning and parent only learning.

6.6.2 Results

6.6.2.1 Visualisation

For each experimental run, three coloured columns are produced. A single row of one column shows the signal used (according to the colour scheme shown on Colour Plate 6.1) by each agent in one generation to signal a particular meaning. Successive generations are shown below each other to form the columns. Thus, column 1 shows the signal used to represent ‘meaning’ 1 by each and every agent over some number of generations. Column 2 shows the signals for meaning 2 and column 3 the signals for meaning 3.

The position of an agent in the spatial array is the position used to plot the signals generated by that agent, as demonstrated in Figure 6.3. With plots such as that shown, the existence of local dialects can be observed across the population, the lifetimes of which can be viewed down the columns.

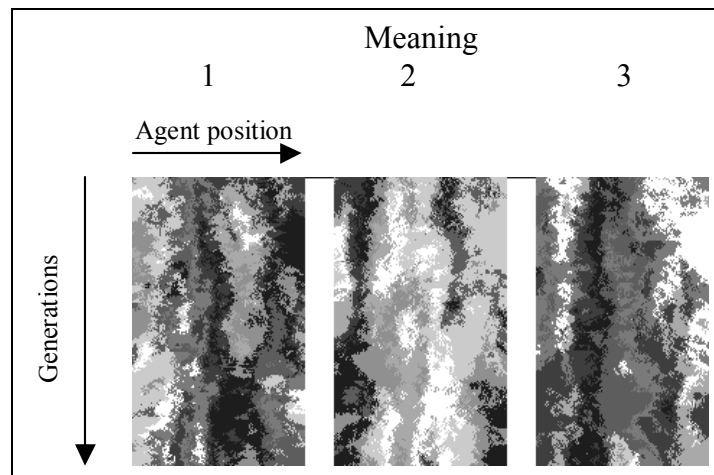


Figure 6.3. Plotting the signal used by each agent in a generation for one of the ‘meanings’ results in a line. Showing successive generations reveals the changing use of signals for the meaning over time. This diagram shows the plots for each of the three meanings, showing the evolution of the communication schemes. Also see Colour Plate 6.2.

6.6.2.2 Experimental Results

Experiments were run under both sets of conditions described in section 6.6.1. Exemplar resultant plots are shown on Colour Plates 6.2(a) and 6.2(b). The diversity obtained is maintained for many thousands of generations – more than can be viewed in a single plot. The results shown in Colour Plate 6.2 each show one thousand generations of language evolution – but the generations shown in every case are the 99,000th to the 100,000th generations.

For the results shown the neighbourhood size is determined by a normal curve of standard deviation 1.5. Other values were used, with similar results – larger standard deviations resulting in dialects covering more of the population. The effect of neighbourhood size is investigated more closely in a following section.

The model had been tested a number of times for 1000, 10,000 and then 100,000 generations without convergence to a single global dialect being observed. Diversity is clearly maintained in the system for long periods, but not infinitely. Running ten test cases for one million generations each, three resulted in communication schemes which had converged across the whole population: for any one of the three meanings, every agent in the population used the same signal as every other agent.

Without any signal mutation, or errors in learning, once the signal schemes have converged, they remain converged. While diversity may be maintained in the model for a long period, it may not survive over many generations, and cannot re-appear once it has

been lost (although later we will look at ways in which signal diversity can emerge in populations with converged signal schemes).

6.6.2.3 Measuring Diversity

The results appear to show that local dialects emerge within a population that does not share a global signalling scheme. This can be investigated using Information Theory.

An entropy (average uncertainty) value, H , can be calculated for each of the meanings for a population. That is, for a given internal state we can determine how much uncertainty exists in the range of signals produced by the agents to signal that meaning. If all the agents produce the same signal then there is no uncertainty. If the signal production is diverse, there is high uncertainty. For a system with eight possible signals the maximum possible uncertainty is 3 bits.

If many local dialects have been formed then the values calculated for H for the signals produced by the whole population for a given meaning should approach the maximum possible value. The more dialects that have formed, the closer to 3 bits the value of H might be expected to be. Taking a localised subgroup of agents from the population, the values for H should be significantly smaller. If the subgroup speaks a common dialect, then there might be very low uncertainty.

Calculating H for each of three internal states, for the final generation, in each of 10 experimental runs of 100,000 generations gives an average of $H = 2.2153$ bits (with a standard deviation of 0.3731). Thus, due to the variation in signal usage over the whole population there is a high uncertainty.

Repeating the calculations, this time for a localised sub-group of the population gives a different picture. The calculations here are based on the signal use of a continuous group of sixteen agents from the centre of the spatial distribution. This time the average uncertainty, H , is 1.1278 bits (with a standard deviation of 0.3744). These sixteen agents have been selected in each case without regard to any possible dialect boundaries – so this lower uncertainty exists even though there may be differences in signal use within these subgroups.

Uncertainty values can also be used to help determine if the diversity is reduced over time. We compare the average value of H after 100,000 generations with that determined after one million. If there is a significant change, then it is likely that diversity has increased or decreased over time. Statistical t-Tests are performed on the sets of H values found under each set of conditions to determine if there is a significant difference between the sets of

results. These results are summarised in Table 6.5. In obtaining the results from the results after one million generations, we use only the data from runs where convergence has not occurred. Including these results shows a significant difference in the level of uncertainty in signal use, but without them there is no significant difference apparent in the two sets of results.

100K vs 1M Generations	Including Converged Data Sets	Excluding Converged Data Sets
One-tailed t-Test	0.002086	0.343076
Two-tailed t-Test	0.004172	0.686152
Paired t-Test	0.002283	0.41212

Table 6.5. t-Test results determine confidence that two data sets may be drawn from the same distribution – whether there is a significant difference between the data sets. Excluding converged results, the uncertainty values (measuring diversity) after one million generations do not significantly deviate from those after 100,000 generations.

Where convergence has not occurred, it is not possible to determine from the uncertainty values whether a data set has been drawn from a population after 100,000 or after one million generations. Except in cases where convergence occurs, there is no appreciable reduction in the signal diversity over time.

6.6.3 The Effect of Neighbourhood Size

The results gathered so far are all for a single neighbourhood size. Intuitively, as the neighbourhood size changes so will the pattern of signal distribution. With smaller neighbourhood size, communication will be more localised and there may be a greater amount of dialect diversity. With larger neighbourhoods, communication will take place over larger areas, and diversity may be reduced.

In the results described so far, diversity is maintained for extended periods of time, and tests with smaller neighbourhood sizes do not cause any change to this result. Of more interest is what happens with the larger neighbourhoods; we investigate whether the population will converge on a single language at a faster rate.

Extending neighbourhood size to infinity results in a single neighbourhood encompassing the entire population, where a uniform random distribution is used to select partners for signal learning and testing. A uniform distribution is used here to ensure that partner agents are selected without regard to distance. Colour Plate 6.3 shows some typical results of signal negotiation under such circumstances (other parameters are unaltered from those given above).

Under these conditions a single global dialect rapidly emerges, yet one that may have some diversity within it.

Plate 6.3(ii) and (iii) show the typical appearance of the resultant communication schemes. The patterns that emerge are similar to patterns of white noise. But for each of the meanings only a subset of the eight possible signals is used. For example, in Plate 6.3(ii), only green (0,1, 0) and turquoise (0, 1, 1) are used to represent meaning 1. Only white and yellow are used to represent meaning 3, and the remaining four colours for meaning 2.

Due to the redundancy in the signalling capability, this allows the agents to correctly interpret the signals that they receive. In effect, a single global dialect has emerged, one that is tolerant of minor variations in signal use between different agents.

Eventually the communication schemes converge, as shown in Plate 6.3(i). This plate shows 1000 generations from the 1000th to the 1999th generation of a sample model run. Convergence happens more rapidly than previously – 40% converge inside 10,000 generations, and all test runs converged by 100,000.

6.6.4 Analysis

These results show that, given spatial limitations on agent interactions during language acquisition, dialects may form. The resultant diversity can be preserved for many generations. Without a source of innovation, the diversity may eventually be lost. So diversity does not last infinitely, but does survive for a very extended duration.

6.6.4.1 Maintaining Diversity

The question to answer is then, how is diversity maintained for so long?

If we consider language to be a cultural system, then the result can be compared to work on cultural diversity, such as (Axelrod, 1997b), who considers how different ‘dialects’ of culture may interact. In Axelrod’s model cultures are defined by a number of features, and the more features cultures have in common then the more likely they are to interact and share their culture – leading to increasing numbers of shared cultural traits. Strong differences result in cultural barriers, which prevent the dissemination of culture. This allows some cultural diversity to be maintained despite increasing homogenisation, and ends in polarization, with few cultures, which share few if any features. However, in our model agents attempt to learn from surrounding agents regardless of the degree of difference in their signalling schemes. Some other mechanism must be responsible for the maintenance of diversity.

Each learner in the population forms their own idiolect after learning from a surrounding mix of idiolects and training examples that is potentially unique for every learner. This allows learners to form different signal schemes from those used by their neighbours, and is one factor in the evolution of linguistic diversity. Another important factor is the interaction of agents who may use similar signals for different meanings – which may result in a learned signalling system being different from both of the original systems, as the learner attempts to resolve the conflict.

An example is shown in Figure 6.4




Figure 6.4. Conflicting training signals may result in agents learning different meaning-signal pairs than presented. In the highlighted region, a signal is learned for a meaning despite not being present in the original training set.

Each agent learns three bi-directional signal-meaning mappings. Redundancy in the signal layer allows multiple signals to be mapped to one meaning (but only one reverse mapping). There are two ways in which this can cause ‘novel’ forms to appear. First, as weights are updated repeatedly as an agent learns, the learning for one mapping may perturb the weights that encode another signal-meaning mapping. This might result in a learner using a different signal for a particular meaning than used by any of the agents in the parent generation. In most cases the conflicts are not solved by one dialect succeeding at the expense of the others.

Alternatively, a learner may only be able to resolve the conflict between two different dialects surrounding it, by learning some mappings not present in either. This may occur when the dialects use the same signal for different meanings. The different mappings mean that each agent has a system of mappings – and changing one value may affect the others.

While the precise workings of this are peculiar to this particular model, this artefact does have parallels to real language diversification, where linguistic innovation may introduce new language features previously lacking in a language community, or where contact between different language communities might have quite unexpected results.

Thus, localised learning is able to maintain diversity, even if it alone is not capable of producing diversity. We have determined that local dialects do indeed exist with the emergent signalling schemes, and that attempting to learn from a mixture of dialects can result in continued diversity. This finding is apparent from the many plots produced, and was measured by an information theoretic analysis of the data. At this point we can also see if the results can be compared more directly to those from linguistic studies in dialectology, and this is done in the following section. Following that, we will review how diversity may emerge from a converged communication system.

6.6.4.2 Human linguistic diversity: A comparison

There exist some qualitative similarities between the spatial organisation of dialects within our model and the geographical organisation of human dialects (as described in Section 6.3). The dialects held by agents within our model form a dialect continuum connecting (at the extremes) dialects that are not mutually intelligible. Figure 6.5 shows the results of testing the success rate for correct signal interpretation by agents, where another agent within the same neighbourhood generates the signal. Almost all agents correctly interpret all signals presented to them.

This is confirmed by comparing the average success of interpreting a local signal, versus the success of interpreting signals that may be generated by any other agent – regardless of distance. Over ten runs, the average success at interpreting signals that originate in the same neighbourhood was found to be high at 98.2%. The average success at interpreting signals which may have originated from anywhere in the population (which will include some number of signals from the same neighbourhood) was much lower at 38.7% (standard deviations of 0.347% and 2.745% respectively).




Figure 6.5: The percentage of communicative successes over the spatially distributed population. Average success is over 95%.

It is also possible to map a correlate of isoglosses, using Hamming distances. The Hamming distance between two binary signals is simply the number of bits different between them – for example (0, 1, 1) has two bits different from (1, 1, 0) and hence has a Hamming distance of two. It is possible to chart the Hamming distances for the signals for each of the three meanings. Figure 6.6 charts the maximum Hamming distance across the three signal pairs when comparing the signals of pairs of adjacent agents across the population – adjacent pairs of agents are selected, and both produce a signal for each of the meanings. Also shown is the result of adding the Hamming distance values for the three signal pairs. The lines marking the change in use of the items do not generally fall together, so boundaries between dialects are generally not distinct. Only at one point does the total Hamming distance total reach 4 (out of a maximum possible of 9, if no signal pairs had any bits in common). This corresponds to the only measured communication success rate score below 90% in Figure 6.5

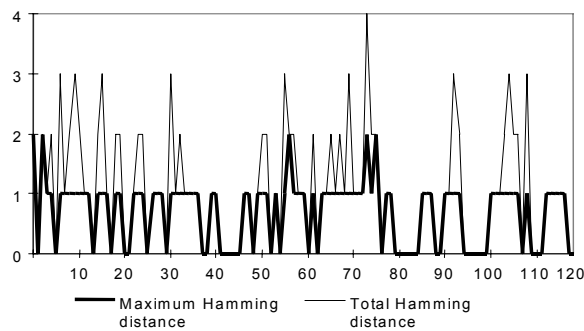


Figure 6.6: The maximum (for any one signal) and total (over three signals) Hamming distances between signals used by adjacent agents in a spatially distributed population.

A more direct qualitative comparison is possible by plotting the accumulated Hamming distance across the population, Figure 6.7. This graph, which shows both areas of linguistic change and regions of dialect stability, is similar in appearance to the schematic shown in Figure 6.1.

These comparisons are useful in showing that the model leads to qualitatively similar results to those observed across the globe. If the model is a reasonable approximation of linguistic transmission in human society, then the processes at work in the model may correlate well with those at work in the evolution of human languages.

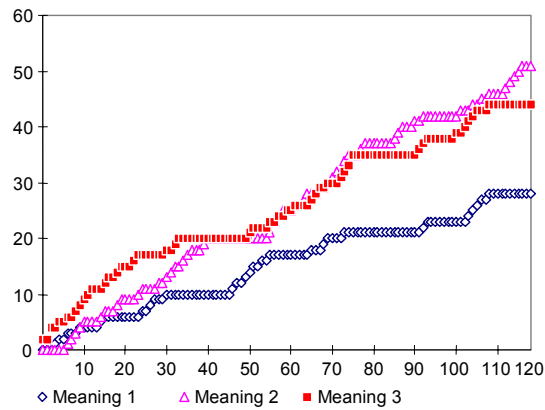


Figure 6.7: The cumulative Hamming distances between the signals used by adjacent agents over the spatially distributed population

6.7 Experiment 2: Diversity from Homogeneity

In the previous experiment, once the communication scheme has converged it remains fixed. Once diversity is lost it is never regained. This compares poorly with human languages, where change and innovation appear to be irrepressible (e.g. Aitchison, 1991; Lass, 1997).

However, there are many potential sources of change and innovation in human language that have no counterpart in our model. The different pressures on language users (see Section 2.5.2) to modify their language use to maximise comprehension and convenience do not exist. Nor does the possibility of errors, which may help drive some language changes, exist. For example, Steels & Kaplan, (1998), demonstrate an ALife model in which errors in signal production, perception and interpretation can cause change in agent lexicons.

Further, grammars are large systems. It is widely recognised that the evidence available to children during grammar acquisition contains contradictions and is insufficiently detailed for error-free acquisition. Language can then be viewed as a system with imperfect replication – in any attempt to replicate an existing language, a number of changes will occur, and the learner will actually acquire a ‘new’ language.

One way to represent this is to add noise to the signals that are presented to learners during training. It is not simple to quantify the incidence of errors and misunderstandings in language use, or to quantify the effect of ambiguous evidence. However, using only small noise values increases confidence that the amount of noise used is not unreasonably inflated.

Another potential source of diversity is individual bias in language learning. If different language learners have innate biases that influence their language acquisition, there is a greater likelihood of learners faced with similar evidence acquiring different languages. As well as unique and individual accents, that allow friends to identify one another by their speech, there may also be subtle variations in individuals' grammars and lexicons.

Individual bias can be modelled by initialising the weights of the language agents to small random values. This may influence learning, preventing agents from converging to the same signal scheme.

With either option it is not necessary to show whether this noise is sufficient to prevent convergence. Rather, it is simply required to demonstrate that noise and/or innate bias can disturb a converged signalling system, returning it to a diverse one. In doing so, it will be demonstrated that even under unusual circumstances, should convergence occur it will not last. This we attempt to do now.

6.7.1 Experimental Setup

The model setup is initially the same as described in Section 6.6.1, with a few differences. Learning examples are provided only by the parent generation – there is no learning from peers. The effect of this is to help to preserve existing systems slightly, and slow down change. We also implement either innate biases or signal noise.

The populations are started with an initially converged language – the first generation is set such that all agents use the same signals. In the results reported here, all agents in the first generation use the signal vector (1, 1, 1), or white, for the first meaning, (0, 0, 0), or black, for the second and (1, 0, 0), or red, for the third. The experiments detailed below were also repeated using populations with initial signal repertoires of red, green and blue (where there is an equal Hamming distance between each signal), with the same qualitative results.

6.7.1.1 Noisy Learning

A variable amount of signal noise is used, set at the beginning of each experimental run. This may cause changes in individual signals, and this may in turn affect the learned communication schemes of the language agents. The noise rate is the chance per signal bit of flipping the bit. For a one percent noise rate each bit of a signal has an independent chance of one in a hundred of being flipped.

Two sets of simulations are performed, both with small noise parameter settings: 0.1% and 0.01%.

6.7.1.2 Innate Bias

All agent weights are initialised to small random values. With the finite, limited, learning period, these weights may prevent the agents from acquiring the same signals with which they are provided as training examples.

Again, two sets of runs are performed. In the first, initial weight values are set to a random value in the range 0.05 to -0.05 . In the second, the range is from 0.005 to -0.005 . In both cases, the initial weight value is drawn from a uniform distribution.

6.7.2 Results

The results for both sets of conditions can be considered together, and are qualitatively the same, and some results are plotted in Colour Plate 6.4. In all cases the innate biases or noise is sufficient to cause divergent dialect evolution, and introduce diversity. From Plate 6.4(c) it is evident that, given sufficient time, the original signal scheme cannot be recovered from the current one, although the rate with which diversity is introduced seems greater when noise is present than when innate biases cause the emergence of diversity.

6.7.3 Discussion

Both innate biases and noisy communication are capable of introducing diversity to the signal schemes used by the agents. The imperfect transmission and acquisition of language is thus the vehicle by which change occurs and diversity is reintroduced in this model.

Here, imperfect learning initially leads agents to learn schemes that, while different, are not incompatible with those of the agents in the parent generation. Inspecting the various graphs plotted on Colour Plate 6.4, the signals first introduced in place of the initial three (white, black and red) have only one bit difference, and are introduced without disturbing the chance of successful signal interpretation. The redundancy of signals allows these changes to take hold, by allowing agents to learn different, but compatible signals.

While human language is much more complex than any simple signalling scheme, let alone one which uses just eight signals to communicate just three meanings, it also contains much redundancy. Lass (1997) argues that ‘linguistic junk’, a consequence of redundant signalling, is a key feature of human language for language change. Our results here agree with Lass.

6.8 Summary

We opened this chapter with observations of what appears to be signal diversity in the negotiated signal schemes created by agents in the previous chapter. We noted that this

spontaneous emergence of diversity seems to contradict some existing arguments on the causes of linguistic diversity, and conducted an additional experiment to investigate the findings.

Having done so, and in order that we might further investigate the emergence of linguistic diversity, we presented a brief review of patterns of linguistic diversity as seen in human society. This was done in order that we might have some basis for comparing the results found in our model against those found in the real world.

We then proceeded to review some work on mathematical models of linguistic diversity. We saw how some of this work leads to conclusions at odds with the real-world evidence. Drawing our own conclusions from a simple model developed by (Cavalli-Sforza and Feldman, 1978), we similarly found an outcome that would be unexpected in the real world.

Then a number of experiments were conducted, under different conditions, to study the emergence and cultural evolution of dialect diversity in the artificial life model previously developed in this thesis. The structure of the agents was fixed for these experiments. This removes from the model those features that attempt to model biological evolution. Instead, by allowing agents to learn their signalling schemes from agents in previous generations, the cultural evolution of language is modelled.

It was found that the patterns of signal diversity found in the model had some significant similarities to the patterns of linguistic diversity found across human languages. Specific results also included observations that the continued existence of diversity could not be guaranteed, although under certain conditions diversity will emerge in populations with initially homogenous signalling schemes. These conditions (noisy learning, or the existence of innate bias) are not unrealistic, and appear to exist in human language learning.

In the next chapter we will present further work that corroborates the conclusions of this chapter, and develop a theoretical framework for these conclusions.

Chapter 7 Cultural Evolution in an Agent Based Model of Emergent Phonology

In the last chapter, we showed how some minor modifications to our model allow us to model the cultural evolution of language, and to obtain results with clear qualitative similarities to the real-world cultural evolution of language.

In this chapter we show that there is a need for additional corroboration, and then develop this corroboration by making use of a different, quite distinct, model. We finish by setting out the significance of our results, and putting forward an argument for the cultural evolution of language *without* linguistic selection.

7.1 Artificial Life and Micro-Simulation Models of Linguistic Diversity

The work described in the previous chapter is not the only work that has been carried out using computational models to investigate the evolution of linguistic diversity. Some of this other work is worth reviewing, starting with work which disagrees with our own.

7.1.1 Functional Requirements for Diversity

In Chapter 2, Sections 2.5.2 and 2.5.6, we reviewed some of Daniel Nettle's arguments on the causes of linguistic diversity. He has also provided backing for his arguments using a variety of simulation-based models.

First, in Nettle and Dunbar (1997), a model is developed which shows how dialects may be used to indicate group membership, and how such a marker may be used in the evolution of cooperation. By using these markers, groups of cooperative agents are able to resist invasion from non-cooperative individuals. This is used as the basis for an argument that dialects emerged for this *reason*. This does not necessarily follow however, and could be another example of exaptation (Section 2.4.2) – applying a novel use (group marking) to a feature (dialect diversity) which has already emerged in the population.

Nettle has presented two further models that support his arguments that social status and social functions of dialect differences are pre-requisites for the emergence of dialect diversity (Nettle, 1999a; Nettle, 1999b).

The model presented in (Nettle, 1999a) arranges language learners into a series of small groups, the language used consisting of a model of a vowel sound system. Learners pass through five life-cycle stages, and all language acquisition occurs during the first stage, where the new language agents learn from the other agents in the same village. Each group

contains four individuals at each of the life-cycle stages (twenty in total at any time). After the fifth life stage, the elderly are replaced by a new set of infants. The infants each learn a sound system according to the set of sound systems in use by the existing group members, plus a small amount of noise. After this, all the individuals are ‘aged’ one stage. No learning occurs after the first stage.

It is shown that, unless the groups are completely isolated from one another, diversity does not emerge. Adding in social status changes the findings significantly. Each individual has a 25% chance of gaining high social status after the first life stage. Learners only learn language from those individuals with high status within the village. Otherwise, for any vowel, the sound learnt is the average formant frequency values used by all of the adults in the population for that vowel, plus a slight perturbation due to noise. With social status included small differences between groups may become magnified over time and it is found that contact between groups no longer eliminates diversity.

The model presented in (Nettle, 1999b) has many major differences, but retains the same agent life-cycle, where agents pass through five life-stages, before being replaced by new learners. Again, learning only occurs during the first life stage. Apart from this there are few similarities. Inspired by social impact theory (Latané, 1981), there are no sub-groups within this model, all agents existing on a single spatial array. Instead of learning vowels, agents acquire one of two grammars, p or q. In determining which grammar a learner acquires, the impact of all the surrounding grammars is calculated. This uses a sum of all the surrounding grammars, weighted by distance. Then, if the result is in favor of one grammar, that is the grammar acquired. Several factors may be varied in this model, but the general finding is that sustained diversity requires that social status exerts a very large influence upon the acquisition of grammar.

These last two models each have design features which lead directly to these results. In the former, vowels are learnt by an explicit averaging of the vowels in use already in the local group. In the latter, the impact measurement and forced selection of a grammar from one of two distinct grammars – without the possibility of acquisition of elements of different grammars – is a form of thresholding. This thresholding forces the grammar learners to acquire the more commonly used grammar variant within their neighbourhood, except where a social-status weighting is introduced.

Nettle argues in his work that the effects of averaging and thresholding would work to stifle diversity, were it not for the effect of social status, and uses these models as demonstrations. He then uses models in which these are enforced by the language

acquisition rules he has built in. It is not proven that under more realistic learning conditions, where language is acquired as the result of many interactions, or where there is a possibility of learning grammars that are different but compatible with surrounding grammars, that averaging or thresholding will prove to be the barrier to diversity that Nettle argues they are.

7.1.2 Other Models of the Evolution of Linguistic Diversity

Arita and Taylor (1996) present what is possibly the first attempt to explain the origin of linguistic diversity using a micro-simulation model. They hold that it is the spatial distribution of individuals that is the key factor in the emergence of dialects. While this is a plausible position, it is not strongly supported by their model, which relies on *genetic* mutation for the emergence of linguistic diversity. Language is inherited, with mutation producing diversity and learning leading to increased convergence. If the spatial distribution of speakers is indeed a factor in the emergence of dialect diversity, then it must be able to work when the only means of language transmission is through learning – as it is for human language

Innate language is again used by Arita and Koyama (1998), in their investigation into the evolutionary dynamics of vocabulary sharing. Mutation rate is again identified as being an important factor in the emergence of diversity in the vocabularies, but without an identified linguistic equivalent. The degree of vocabulary sharing is also related to the availability of resources. Rather than vocabularies, it is cooperative strategies that are being evolved here, as evidenced by cases where the evolved communication strategy is not to communicate at all. Accordingly these last two models give few insights into the causes of, and influences on, human linguistic diversity.

7.1.3 Related Models

In some cases, it may be that models that have been developed to study other systems may, in some way, be relevant to the study of linguistic diversity.

For example, Axelrod (1997b) presents a model to investigate the dissemination of culture through a spatially arranged population. In the model, neighbouring sites may interact if they have at least one cultural trait in common, and as sites interact they share traits and slowly converge. Eventually a stable distribution emerges where a limited number of groups survive, within each group all sites having identical sets of traits, and no traits being shared with sites belonging to neighbouring groups. Viewing language as a cultural trait, this is obviously relevant to the evolution of dialects. The results are at odds with

observed phenomenon in human language, however (see Figure 6.2, and related discussion).

In other cases the relevance of the results to the question of the origins of linguistic diversity is clearer. For example, the previously mentioned work of Kirby (1998), section 2.4.3, builds a simulation model to show that Universal Grammar constraints may not be innate constraints at all, but merely the outcome of learning over time leading to a reduced set of surviving grammars. In this work Kirby shows results indicating the existence of geographically distributed dialects of grammars.

Following on from (Kirby, 2000), which demonstrates the emergence – through cultural transmission – of compositional language in a population of language learners, (Brighton and Kirby, 2001) present a study of the eol exploring the conditions under which compositional languages are maintained.

Maeda et al. (1997) examine the effect of language contact. The results show language reorganization after contact is made between populations using different languages. But after the reorganization, dialect diversity is completely absent from the population – again, not results that compare well with those observed in the real world.

Other related work has looked at the process of language change, apart from the question of dialect. Steels and Kaplan (1998) demonstrate how various linguistic and extra-linguistic errors can lead to continued language change. While natural language errors are somewhat more systematic than the random errors introduced in this model, the model successfully demonstrates the large influence such errors may have on language innovation. A similar model, based on artificial neural networks, is presented by Dircks and Stoness (1999). In this model, it is found that noise is not required to maintain competition between forms. This appears to be due to the networks learning one of two similar signals for particular internal meanings – an ability not present in the Steels and Kaplan model, where similarities in the lexical forms are ignored by the agents.

7.2 Dialect in an Agent Based Model of Emergent Phonology

The preceding review has shown that the work of the previous chapter is not the only work using computational models claiming to illustrate the processes which give rise to dialect diversity. Some of these models produce results which appear to directly contradict our own conclusions.

Nettle has produced a variety of different models, all of which support his theory that social motivation is required to produce and maintain linguistic diversity (Nettle, 1999a;

Nettle, 1999b). We have argued that the details of these models are such that they explicitly implement the averaging and/or threshold problems. Human language learning is not based on assessing what the average signal is, or on calculations to determine what signal is used by the majority of people and so explicitly including such rules in a model may invalidate the results of the model.

Yet part of the strength of Nettle's argument is that he has demonstrated the same requirement for social motivation in a number of unrelated models. The model presented in the previous chapter does provide contrary evidence to Nettle's, but additional supporting evidence from an unrelated model would provide useful support to our argument.

With such a goal in mind, we re-implemented Bart de Boer's model of an emergent phonology (de Boer, 1997), extending it slightly in order to investigate the emergence of dialect diversity.

7.2.1 de Boer's Model of Emergent Phonology

De Boer's model is one in which populations of agents form an emergent vowel sound system, through repeated interactions and learning. An earlier model of emergent vowel sound systems has also been presented by (Lindblom, 1998), but this uses more traditional techniques to find optimal distributions of vowels within a vowel sound space without modelling the language users or their interactions in any way. De Boer's results compare well with these earlier results, with the advantage of more closely modelling the processes involved when populations of speakers try to match the sounds that they hear being spoken.

De Boer's model has been thoroughly documented, enabling an independent re-implementation to be attempted (de Boer, 1997; de Boer and Vogt, 1999; de Boer, 2000). This was done, and testing determined that it performed qualitatively the same as de Boer's own implementation.

The agents produce and attempt to learn vowel sounds. Vowel production is based on a mathematical model of the human articulatory system, where tongue position determines the formant frequencies of the generated vowel. Agents use a model auditory system to try to perceive the vowel, and can then try to mimic the sound.

Each agent maintains a list of vowel prototypes. Upon hearing a vowel, a learner agent determines which vowel in the list is the most likely one that they may have heard. This

sound is then produced. The original signaller listens and determines whether the vowel it hears back is sufficiently similar to the original vowel, and with a yes/no feedback signal. Based on the feedback, the learner either updates the prototype for an existing vowel, to shift it closer to the perceived one, or creates a completely new vowel prototype. These interactions are repeated, with learners and tutors randomly selected from the population. Over time, vowel prototypes are added and occasionally merged or pruned until each agent has a set of vowels. With varying amounts of noise, vowel systems with few or many vowels may emerge.

The results of this emergent system have been compared extensively with human vowel systems, and this has confirmed that the vowel systems so produced are realistic and highly possible.

The vowel systems may be viewed by having all of the agents produce the sounds for each of their vowel prototypes and charting the vowels. The vowels generated are represented by four formant frequencies, F1, F2, F3 and F4, but for purposes of display the last three are combined (see de Boer, 2000, p.182, equation 5) into a single value, F2'. Figure 7.1 shows a typical emergent vowel system, from a simulation with a population of 20 agents, and 5000 learning interactions. Added to the chart are approximate positions for some of the major vowel sounds of the English language. Such charts of 'vowel space', plotting vowel sounds according for first and second formant values are commonly used in phonetic studies (e.g. Johnson, 1997, p.105). More details of the inner workings of de Boer's model are included in Table 7.1.

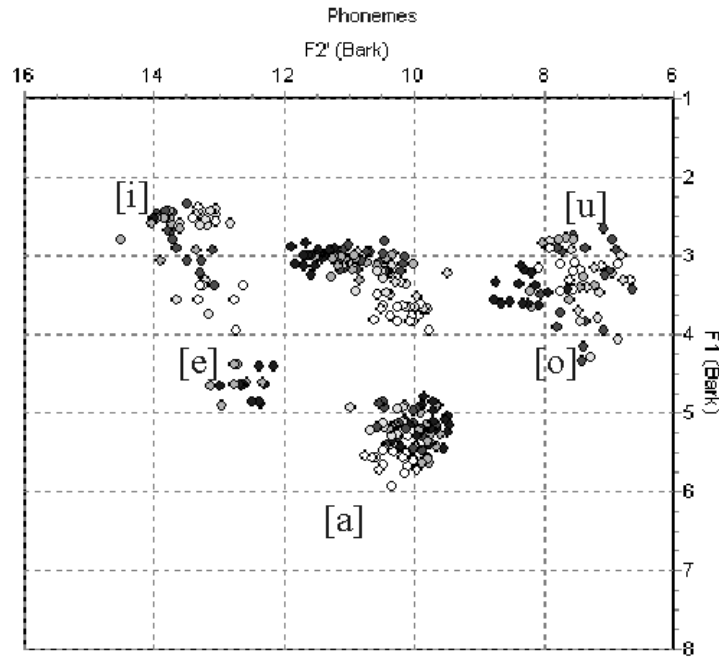


Figure 7.1. An emergent phonology. Clusters appear in areas of the phoneme-space where multiple agents have learned shared vowel sounds. The approximate positions of some the major vowels of English have been superimposed on the graph.

Initiator (teacher)	Imitator (learner)
If $V = \emptyset$ then Add random vowel to V	
Pick random v from V , Increment count of uses of v Produce signal A_1 from v	
	Receive A_1 If $V = \emptyset$ then $v_{new} = \text{Find phoneme}(A_1)$ $V = V \cup v_{new}$ Calculate $v_{received}$ Produce signal A_2 from $v_{received}$
Receive A_2 Calculate $v_{received}$ If $v = v_{received}$ then Send non-verbal feedback: success Increment count of successful use of v Else Send non-verbal feedback: failure	
Do other updates of V (see below)	Receive non-verbal feedback Update V according to feedback signal Do other updates of V (see below)

Table 7.1. Extract from the basic rules for agent interaction in de Boer’s model of emergent phonology (from de Boer, 2000). V is the set of vowels possessed by an agent, and v is some vowel. Signals A_1 and A_2 are the articulations of the selected vowels.

During the ‘other updates of V ’ step of de Boer’s model, bad vowels (which do not correspond well to vowels heard) are removed, vowels which are close to one another are merged and, with a small probability, new vowels are randomly added to V .

7.2.2 Experimental Setup

The model was enhanced to include a larger population of agents, spread across a spatial array. 100 agents were arranged in a single, non-toroidal, line. The algorithm for agent learning is not altered, other than to enable the neighbourhood-based selection of partners. Once an agent has been selected a partner is required for learning. The partner is selected from a position along the line on either side of the original, within a limit of ten agents distance. Agents near the ends of the line have as a consequence fewer other agents with which to communicate. The selection of partners is from a uniform distribution.

Table 7.2 sets out the parameter settings for the simulation detailed in the following results. As before, a number of simulations were performed for a variety of parameter settings, some of these are discussed below.

Parameters			
Population	100	Training rounds	25,000
Neighbourhood size, d (+/- d)	10		
Noise (%)	15		

Table 7.2. Parameter settings for emergent phonology model

The large number of training rounds ensures that each and every agent will receive of the order of 200 training examples (the number of training rounds being determined by the population x 250). Varying the number of training rounds thus allowed tests with larger and smaller populations to be performed, the results of these being supportive of the assertions following.

Neighbourhood size was determined by the size of the population divided by ten. This figure ensures that there is reasonable distance between the far ends of the population. The actual neighbourhood size is double this, as agents can communicate with others on either side up to distance d .

In de Boer’s thesis the effect of varying noise is well documented. Smaller noise values allow the emergence of vowels systems with more vowels than occur with larger noise values. The value used in these experiments is in the mid-range of values used by de Boer. This would lead us to expect vowel systems of around four to six distinct vowels to emerge.

7.2.3 Results

Running the model produces the output shown in Figure 7.2.

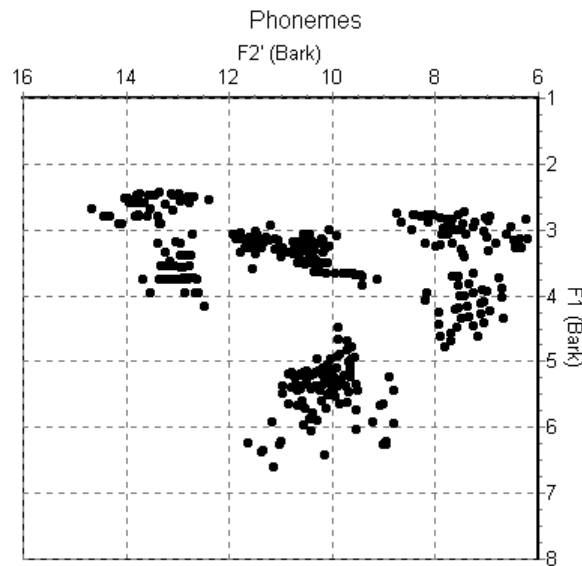


Figure 7.2. The emergent vowel system of the population. Depending on how the vowels are categorized, there appear to be around six distinct vowels in use.

It is not clear in Figure 7.2 how many distinct vowels have emerged – the precise answer depending on how the individual sounds are categorized. Possibly the clearest categorization would be to count two vowels at the top right of the chart, two at the top left, one top centre, and one more at the bottom. This gives a total of six vowels, within the expected range. At the bottom of chart, there are some breakaway vowels. A more detailed examination reveals what is really happening.

The diagrams in Figure 7.3 show the same emergent vowel system. The population has been split into five arbitrary contiguous groups. The first twenty agents of the population form the first group, the next twenty are placed in group two, and so on.

Some of the individual diagrams remain a little unclear – it is not always obvious whether one or other of the clusters represents one or two vowels. Even within a single group, there is some spatial distance, and it is possible that different agents within a group have learned slightly different sets of vowels. The groups themselves have not been chosen with regard to how close the vowel systems of the individuals within the group are – rather, the groups are a completely arbitrary division of the population according to spatial position. As such it should not be expected that the vowel systems will be extremely close within these groups, with distinct differences only occurring between adjacent groups.

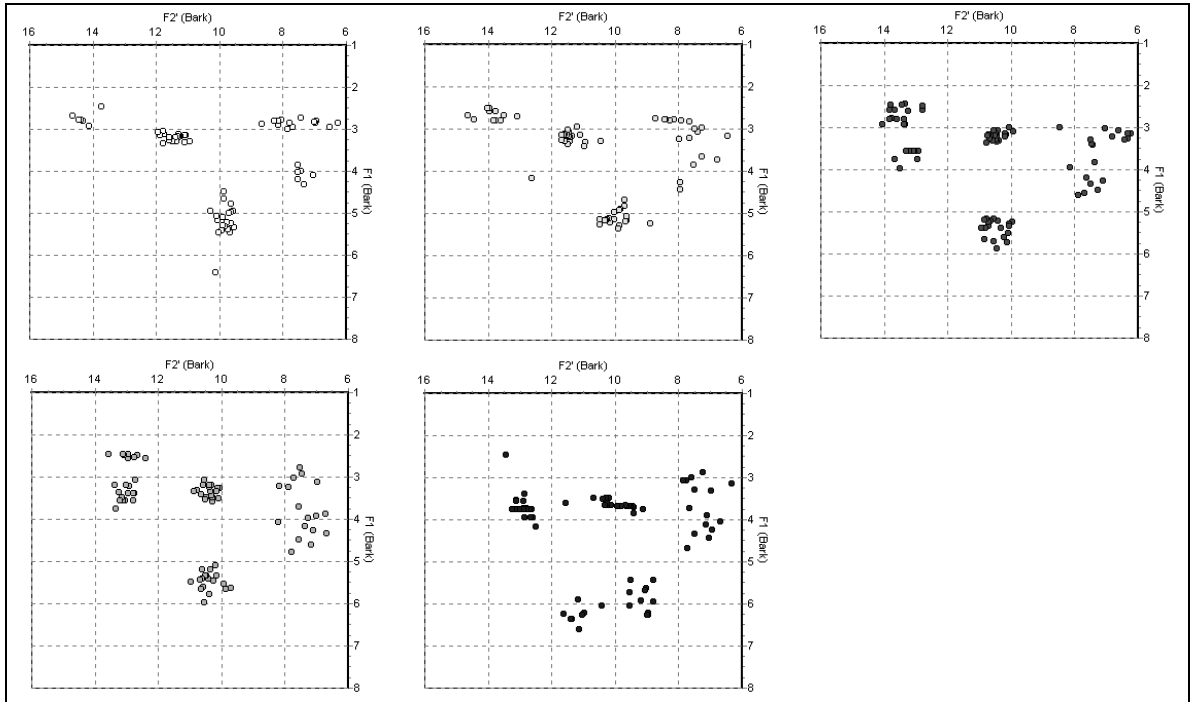


Figure 7.3. The emergent vowel system of the population. Each diagram shows the phonemes used by a different contiguous sub-group of the population. Top row shows groups 1-3, bottom row groups 4 and 5.

However, it would appear that most groups have developed a four or five vowel system, which is largely shared amongst the agents within a group. Figure 7.4 emphasises the differences that exist across the population. In this figure, the phonemes of the first and the last groups are shown together (as white and black dots respectively).




Figure 7.4. Emergent vowel systems of the first and last sub-populations.

As with the previous work (sections 6.6 and 6.7), a dialect continuum has emerged in the population. Minor changes exist within neighbourhoods, allowing successful communication therein. Across the population, more major shifts and differences exist.

Although the model is entirely unrelated the qualitative result is the same: the negotiation of a communication scheme/phonology within a population, where neighbourhoods limit interactions, gives rise to emergent dialects without any requirement for any need or motivation to create the different dialects.

The de Boer model has some advantages over our own for studies of language change and diversity. There is no fixed number of signals, or of ‘meanings’, leading to systems which are more open, more like those found in human language. The differences in vowel use by different sub-populations, or over time (de Boer, 2000, p.190-192) generate what appear to be *chains* of changes. The close parallel of the emergent sound systems to those of human language makes comparison of results to observed changes in human languages possible. As such, there is potential to use this model to further study the development of push/drag chains (systems of change where movements in clusters create spaces for other clusters to move into, creating a ‘drag’ effect, or where the movement of one sound forces another to move apart, a ‘push’ effect’, (King, 1969)) in vowel sound systems.

7.3 Discussion: Towards a Modified-Neutral Theory of Language Change and Diversity

The experiments in this chapter and the last, both demonstrate a form of cultural evolution without any cultural selective pressure. Thus these experiments demonstrate one form of neutral evolution of language – the evolution of language without selective pressure on adopted forms being exerted by social bonds or factors.

7.3.1 Neutral Evolution Revisited

Neutral evolution, and its application to linguistic evolution was discussed previously, in section 2.5.6. In light of these most recent results, we look again at neutral evolution, and how it may apply to the eol.

Neutral evolution is evolution in which selective processes do not operate (Kimura, 1983). Neutral evolution, also known as genetic drift, can occur anywhere where there is variation in a population but where none of the variants has any specific selective advantage. By chance, rather than by adaptive selection, one variant may become a new norm. The resultant species or languages will be distinguishable from their ancestors. For biological entities the definition of ‘neutral’ change is quite clear, as being any genetic mutation which does not affect the reproductive success of the creature. In the case of language there are a number of possible interpretations.

The first is an analogue of the original meaning. A neutral change in language may refer to some linguistic change which does not affect the ease, or difficulty, of acquiring that language. Christiansen (2000) and Kirby (1999) have looked at how the ease with which a particular language variant can be learned can influence its survival, showing that hard to learn variants are selected against and may not survive. What is left may then be a large number of variations that are all – possibly approximately – as easy to learn as each other. Selection amongst these variants may then proceed in the manner of neutral evolution. This will henceforth be referred to as *linguistically neutral* evolution.

A second interpretation, closer to the original biological meaning, is that an evolutionary neutral linguistic change is one which does not affect the reproductive success of the language users. There are a number of ways in which language influences evolutionary fitness, and determining how speaking a particular language sways the speakers' reproductive chances is not entirely straightforward. There are two distinct ways in which the fitness of language users could be affected by changes in their speech, according to whether the changes affect the communicative success or effectiveness of dialogues or if the changes affect their social standing within their group. Changes that do not affect the success of communication could again be considered as being linguistically neutral. Changes that do not affect the fitness of language users according to social position will be referred to as *socially neutral* evolution.

7.3.2 Adaptive, Maladaptive and Neutral Change

Where changes affect the communicative success, whether the benefits are to be gained by communicating ideas, observations or by gossiping (Dunbar, 1996), adaptive changes will be those which make communication, both signalling and interpreting, easier.

In arguing that language change cannot be socially neutral, Milroy (1992) argues that from the perspective of communicative fitness, all change is dysfunctional – by not speaking precisely the same language as others an increased chance of misinterpretation is introduced (also see section 2.5.2). This is used to justify arguments that change *must* be socially adaptive. Such a view is extensively denied by Lass (1997). Lass notes that language features a high degree of redundancy and as a result it is possible to make quite significant changes before communication is adversely affected, a point also made by Pinker (1994). While some changes may affect successful communication, there is much room for changes which do not degrade communication.

However, the social benefits of speaking the correct dialect are well-documented (see, for example, Chambers, 1995), and there have been many numerous studies of how dialect markers are used in determining membership of all manner of groups from street gangs to business elite. Social marking may also influence reproductive fitness of language users, as a consequence of increasing cooperation amongst those with similar dialect (Nettle and Dunbar, 1997).

Nettle (1999a) further claims that, with a neutral model, it is difficult to account for diversification without there being complete geographical isolation between groups. Thus, Nettle proposes that in order for linguistic evolution to occur without geographical isolation additional mechanisms are required. Nettle argues that the social functions of language are *required* for the emergence of linguistic diversity, a view shared by Dunbar:

“... dialects arose as an attempt to control the depredations of those who would exploit peoples’ natural cooperativeness”

(Dunbar, 1996, p169)

However, Nettle’s arguments rely on a particular learning model – in which the learners sample the speech of the population and learn an ‘average’ of the language around them. This further relies on the equal distribution of individuals, with a uniform likelihood of any one individual interacting with any other. As recognised by Cavalli-Sforza and Feldman (1978), in any group the amount of influence exerted on any one individual by any one of the others will vary according to a number of factors. This reduces the effect of averaging, and increases the potential for sub-populations to vary from the mean. The different social networks within groups reduce the need for geographical isolation to produce linguistic diversity.

Nettle does not consider the effects of sub-groups within communities, which I suggest is an important feature in the development of linguistic diversity. Dunbar (1996) puts forward 150 members as being a natural group size for human communities, and points to “sympathy group” sizes of between 10 and 15 people. Any change which does occur must be propagated over a series of interactions between individuals. With the existence of groups-within-groups, there is no need for isolation before linguistic diversity can emerge. Further, the averaging effect itself is questionable. For example, for random variation in the formant frequencies of phonemes it may not be the case that such variation will ‘cancel out’, or that the average values will be learned. Phonemic and articulatory constraints (see Lindblom, 1998; de Boer and Vogt, 1999) may prevent the ‘cancelling

out’, and the learned phonemes may be only tolerably close to those heard. Given two distinct forms of a linguistic feature, a learner need not choose just one to learn or to learn some composite, but may learn both. One form may be preferred, but both may be used in varying amounts (such variations in language use often being the focus of studies in sociolinguistics). As well as applying to the lexicon, it has been proposed that language learners learn multiple grammars, so as to cope with the variation in grammars in use around them (Kroch, 1989).

This objection about averaging is extended to Nettle’s computational model in which averaging is explicitly performed. In contrast, the results presented in this and the previous chapter successfully show the emergence of dialects as a consequence of localised variations and interactions *without* any social significance or utility of particular dialect forms.

While particular linguistic forms may mark group membership, other linguistic features and forms might be free of any information that marks the speaker as belonging to a particular group. Where they exist, changes to non-marking features – be they phonological, lexical or grammatical – might be possible without any socially adaptive consequences for the speakers. Indeed, in Milroy and Milroy (1993) such changes are noted amongst the men and women of two Belfast communities, where the authors state,

“it is the group for whom the vowel has less significance as a network marker which seems to be leading the linguistic change”.

(Milroy and Milroy, 1993)

(reprinted in Trudgill and Cheshire, 1998, p192)

By the arguments of Milroy (1992), change must be socially adaptive to overcome the linguistic pressures against change, yet here we have evidence of change without social advantage.

So we have seen that there is evidence and argument to support the existence of both socially and linguistically neutral change in language. Naturally, a change which is socially neutral may or may not be linguistically neutral, and vice versa. Indeed, it is possible that should the benefit in one domain (social or linguistic) of a particular change be strong enough, change may occur even though it is somewhat maladaptive in the other domain. In the next section we look to other theories in linguistics and assess their impact on a neutral theory of language change.

7.3.3 Universality and Uniformitarianism

Two popular ideas about language also lend some support to the idea of language change as being the result of neutral evolution – Linguistic Universality and Uniformitarianism.

7.3.3.1 Universality

A prevalent belief about language in modern linguistics is that no language is superior to any other. All languages are not the same, but they are all broadly equivalent in their overall expressive and communicative powers. While particular concepts may be somewhat easier or harder to express in one language than in another, on balance it is not possible to say that any one language is ‘better’, and many examples which were previously held to show such differences have since been reassessed and shown to be no more than myth (e.g. Pinker, 1994).

If this is the case, then it would appear to be the case that language change must be linguistically neutral. If it were not then we should expect to be able to find languages that *are* better than others. Linguistic evolution occurs at a significantly faster rate than biological evolution, and over the millennia, as different languages evolved at slightly different rates, some examples of more primitive languages would be left. But as languages round the world have been studied it has been noted that while there are many primitive cultures, there are no primitive languages. Universality, however, says nothing on whether language change is socially neutral.

7.3.3.2 Uniformitarianism

The Uniformitarian Principle as applied to linguistics implies that all languages are subject to a common and unchanging set of rules (e.g. Lass, 1997, p. 26-32). Over the millennia in which languages have been evolving, the same processes and the same rules have always applied, and the same rules and processes apply to all languages in the world. Such a principle is seemingly supportive of – or at least, not opposed to – the idea of neutral evolution. This does not imply that all language change is socially neutral, but we could suppose that were *all* change socially driven that changes in society would have more marked effect on the rules and process of language change. Thus, the uniformitarian principle gives some support to the idea of language change being the result of socially neutral processes.

7.3.4 Relativity and Non-Uniformitarianism

So, from the above it seems that some of the prevalent views on language change are either supportive of, or not opposed to, a theory of language change through neutral evolution. However, there are those who continue to argue against, and to present new arguments against, linguistic universality and the uniformitarian principle. If our theory of neutral evolution is to be robust, then it should be shown how it is not totally incompatible with the counter positions to universality and uniformitarianism. In this section we highlight some of the possible challenges, and address them in the following.

7.3.4.1 Linguistic Relativity Hypothesis

The idea of linguistic universality is often defined in opposition to the notion of linguistic relativity. The linguistic relativity hypothesis holds, first, that speakers of particular languages have their thoughts shaped by the languages they speak, and second, that different languages may be significantly better or worse than others at expressing different concepts (Lucy, 1999). This is also commonly known as the Sapir-Whorf hypothesis, after the works of Edward Sapir (Sapir, 1949) and Benjamin Whorf (Whorf, 1956), and also as Linguistic Determinism.

While highly debated, and currently out of favour (Pinker, 1994), the case for dismissing the linguistic relativity hypothesis is not proven, and it still has some prominent supporters (for example, Lucy, 1999; Slobin, 2003).

We can suppose that where significant differences exist in how two different languages represent concepts, one of the languages may be better adapted to a particular environment, society or task in general. The principles of the linguistic relativity hypothesis would therefore hold that some of the changes that have occurred in language histories have been linguistically adaptive, bestowing adaptive benefits on the speakers of the new variants.

If languages can be adapted to particular cultures, then must the uniformitarian principle also be reassessed? This is discussed in the following paragraphs.

7.3.4.2 Non-Uniformitarianism

While uniformitarian principles are held by many in linguistics, at least broadly so, (Newmeyer, 2000) reminds us that this uniformitarian assumption is not a safe one. Newmeyer argues for two possible types of non-uniformitarianism, ‘non-U’. Strong-non-U supposes that there are functional forces acting on language that are somehow culturally

determined. There is evidence that some languages are adapted towards the functional needs of particular societies, and so this is possible.

Weak-non-U supposes that there are constant functional forces acting on language, resulting in particular direction to linguistic evolution, such as a general tendency to change from OV order grammars to VO grammars. Thus, in Weak-non-U, the rules and processes are constant but they effect continued directed change on languages over time. Other arguments that the evolution of languages is in fact directed, and not neutral, have been presented in (Bichakjian, 1999).

In both cases it is implied that over time a selective pressure is being applied to languages, and hence change is the result of adaptive and not neutral change. In the strong case, the pressures applied to languages are themselves the product of the societies using them, leading to “non-accidental correlation between ‘purely’ grammatical features and aspects of culture, climate, and so on.” (Newmeyer, 2000, p.166).

7.3.4.3 Directed Change in Sound Systems

The discovery of directed changes in language evolution is not new. As noted in Chapter 2, much of historical linguistics takes a structuralist view where language is considered a system, and where changes in part of that system may lead to changes in another part. Grimm’s and other sound laws are based on observations of such change (Adamska-Sałaciak, 1997). These sound laws represent other examples of directed change – examples which are replicated in different languages across the world.

This is the basis for one of the arguments against neutral change proposed by (Nettle, 1999a), who argues that a neutral model should result in a random pattern of linguistic diversification, whereas observed patterns of change have structural correlations, representing parallel evolution. Again it would appear that directed change in sound systems must be linguistically adaptive.

7.3.5 Neutral Change in a Relativistic, Non-Uniformitarian World

Linguistic relativity and non-uniformitarian ideas, as well as directed change in sound systems, appear to present strong challenges to the idea that language change may be an evolutionarily neutral process. They are, however, challenges that can be answered.

7.3.5.1 Neutral Change In Sound Systems

There is currently no evidence, and few arguments, that particular sound changes have *benefited* particular languages. Based on current knowledge, it is not possible to say that

the English language was in any way better after the Great Vowel Shift (Baugh and Cable, 1978) than it was before – despite the many changes to the vowel sound system that occurred.

If the result of many individual changes is a sound system with no obvious linguistic benefits over the sound system that existed before the changes, then the result of the sequence of changes must be adaptively neutral. In a sense it is the individual sounds which are co-evolving over time, causing changes to the sound systems they form. The directed changes that appear are a consequence of pressures on each sound in a system from the others. While the changes are adaptive from the point of view of the survival and evolution of individual sounds, they are not adaptive changes for the speakers.

Accordingly, we see that directed changes can still be evolutionarily neutral. This thought experiment shows an adaptive change in one domain (sound) as being neutral in another (human survival). But beyond this, could language change – directed or direction-free – be an example of neutral evolution in its own domain?

The pressure on language users to use vowels which are distinctive enough to be readily distinguishable limits the space of possible changes, selecting against ‘bad’ changes. If a change occurs selecting an adaptively equivalent variant, that change is adaptively neutral. In any sound system, there may only be a small number of such possible changes. If all selections were between adaptively neutral variants then, over time, it would most likely be noted that from similar start points, similar changes would occur in different sound systems.

This notion that there may be a set of variants within which selection can occur without adaptive benefit is known in biology as a *neutral network*.

7.3.5.2 Neutral Networks in Language Evolution

A neutral network (Huynen, 1995) is the term given to a set of variants of a form which are of equal fitness. The idea of neutral networks can be applied to linguistics at the phonemic, lexical and grammatical levels. In each case, there may be a set of possible changes which do not affect the linguistic ‘fitness’. Selection can occur freely within these limited sets. By having a constrained choice of selectively neutral changes, neutral does not equate to a purely random process. The existence of many other forms that are selected against does not mean that the selection that *does* occur is not evolutionary neutral.

7.3.5.3 Neutral and Adaptive Evolution

Our final argument contends that it does not even matter if it is shown that many examples of language change *are* adaptive. There is still good reason to believe that neutral change occurs, and is a significant process in the eol. To comprehensively deny neutral evolution, it must be shown that *all* changes are functional. For neutral evolution to be considered a factor in linguistic evolution it only has to be shown that *some* changes are not functional, a point also appreciated by Lass (1997, page 354).

As noted above, the concept of neutral evolution comes from biology where there is certainly no lack of examples of adaptive change. It has even been claimed that the majority of mutations are selectively neutral (Nimwegen et al., 1999). Neutral evolution encompasses all of the changes that do not affect fitness. With a high degree of redundancy present in language, there is a great deal of room for such changes.

7.4 A Modified-Neutral Theory of Linguistic Evolution

To summarise, we have argued that no functional or adaptive benefits are required to create linguistic diversity and that diversity should arise naturally from the imperfect transmission of language from users to learners. This represents a neutral theory of linguistic evolution and we have shown that this could well be responsible for diversity in language dialects.

We have reviewed some of the objections to a neutral theory and shown them to be unconvincing. Accordingly, social or linguistic functions are seen to be unnecessary for the emergence of diversity. What then is the role of social and personal motivation in language change? To say that adaptive benefits are not required for the evolution of diversity is not to say that such benefits do not exist, or that they do not influence the evolution of languages. Indeed, classic studies such as that of language change in Martha's Vineyard (Labov, 1972) show that social factors do exert a strong influence.

Accepting that language changes are influenced by social pressures on language users, we can question why language users adapt their language according to such pressures. Is there something remarkable in the human ability to determine significant social information simply from accent and dialect, without regard to the content of the speech?

Rather than claim that it is the usefulness of dialect as a social marker that led to the evolution of linguistic diversity, we would argue that the reverse is more likely – that the

unavoidable linguistic diversity has led to increased awareness, and utilisation, of such differences.

We conclude that the neutral evolution of languages is unavoidable, but that it is not the only cause of change – hence, our use of the term a *modified*-neutral theory of linguistic evolution. While neither social nor linguistic function are required to create linguistic diversity – geographical spread and imperfect transmission alone are sufficient – both remain as important factors in the evolution of languages.

7.5 Conclusions

In this chapter we performed a further, short, literature review – this time of work detailing the results of, and conclusions from, a variety of computational models of linguistic diversity or related systems. Seeing that there was some need for additional corroboration of our previous results, such corroboration was provided by performing additional experiments using the emergent phonology model of de Boer. This model generated results which are qualitatively comparable to observations of human language diversity, as did the results of our previous model (Chapter 6).

From the results of this work, and from our reading of relevant literature, we derived our modified-neutral theory of linguistic evolution. This theory holds that particular language changes could not only be *socially* adaptive or mal-adaptive for speakers, or *linguistically* adaptive or mal-adaptive, but that they can be adaptively neutral. Language change and dialect diversity can emerge without any adaptive benefit: such change and diversity being a consequence of the repeated learning of language by different individuals distributed by spatial or social constraints.

In the final chapter, we review the importance of this work – and of the other work detailed in this thesis – and point to ways in which this work could be extended in the future.

Chapter 8 Conclusions

In this thesis, we have reviewed some of the growing body of work which uses artificial life methods to investigate the evolution of language and languages, and presented some of our own artificial life based investigations. In this, we have encountered some limitations of the artificial life method for building explanations of the EoL, but have also found results which we believe should be of interest to the wider community interested in the evolution of language.

8.1 Artificial Life and the Evolution of Language

In this thesis we have seen how computational models can illustrate and illuminate processes that are otherwise hard or impossible to observe, and how they can be useful in improving our understanding of such processes. Thus artificial models can be useful where it is not possible to make direct observations of the phenomenon of interest, such as is the case with the evolution of language.

Despite this, we have argued that such models have limitations in their power in adding to the considerable base of knowledge that exists on the evolution of human language ability. The principal theories that already exist have been developed without the use of models, and to date computational models have had limited worth in advancing any one theory of the origin of language over any other.

There is one key reason for this limitation. The problem domain is very open ended while the models that have been developed have, of necessity, been much more limited. The problem is similar to that discussed by Tim Taylor (Taylor, 1999) on the use of models of self-replicating systems to investigate the emergence and evolution of life and the development of species and organisms of increasing complexity. For the models to work very significant constraints are imposed on the representation of individuals or on the environmental and reproductive processes. By providing set rules and mechanisms for reproduction, or by providing a grossly limited set of possible interactions with the environment, the model will in effect have greatly limited the search space and set of possible outcomes.

This is most evident in some artificial life work on the EoL where the investigator sets out to determine if some application of language might provide sufficient adaptive benefit to explain the emergence of language. A model is then built in which agents

are able to evolve and interact. Typically, however, both the evolutionary changes available *and* the interactions possible are tightly bound, and very limited. Ideally agents should not possess the appropriate mental structures and/or rules for language use, as early hominids presumably did not, but these should evolve. Instead, the rules for language use, and the representational capacity for language use are usually explicitly present from the beginning.

With only one application of language possible, and suitable fitness rewards for agents that are able to use the model language, the experimenter then runs the simulation and waits for random mutation to find the correct solution. This has led to a range of works that each support one of a variety of uses of language as being the prime reason for its emergence (see reviews in Chapters 2, 3 and 5). From an open-ended problem domain, limited models are built that can only support a single theory. Against this background of numerous positive results is the knowledge that, in nature, only one species has evolved language, and only a very few species have naturally acquired learned systems of communication at all. If each of the many posited reasons for language's genesis was, by itself, sufficient cause, then we might expect it to be somewhat more common than it is.

As methods and models improve it may be possible to relax these constraints. But it is clear that the complexity of the EoL is such that it will be some time before simulation models of the evolution of language will be able to significantly add to our knowledge of why and how language evolved. The best we can currently achieve with such models is to illustrate particular points or ideas.

8.2 Methodological Approaches

However, using some of the reviewed methodological principles for conducting artificial life experimentation (reviewed in Chapter 3) can help researchers in finding productive topics to pursue. In particular, work that tries to resolve current questions by building models, where competing theories can be tested against one another, can provide useful insight and useful additional evidence for resolving debate. An example of this approach is in the more detailed findings and arguments that have been drawn from the work of Chapters 6 and 7 compared to that of Chapter 5. Further examples are provided in some of the other work that has been cited in this thesis.

In contrast to work that attempts to explain the origins of language, modelling the cultural evolution of languages appears to suffer less from the problems mentioned

above. The evolution of languages occurs in a system with a predetermined biological substrate. The structure of individual agents, and the basic rules that they follow, do not vary over time, but are predetermined. This is an important difference, and one that bypasses the problems noted by Taylor (1999) about the use of artificial life for modelling evolution.

In our own work, we have tried as much as possible to follow the principles of good practice discussed in Chapter 3. The model developed in Chapter 4 and used in experiments in the following two chapters was a simplified version of an earlier model. In simplifying the earlier model, we were attempting to build a minimal model that was as simple as possible, yet not so simple that it was no longer relevant to the problem being modelled.

We have presented arguments validating the design decisions taking in building the models – such as when deciding what characteristic features of language should be implemented in our minimal model (Chapter 4), or when choosing to build a model of the EoL in which all agents are equally well able to learn how to interpret signals, but not how to produce them (Chapter 5).

We have been equally attentive to the task of verification. In our work on the EoL, in Chapter 5, this verification resulted in our conclusion that our model was unable to provide evidence to favour either the continuous or the discontinuous EoL over the other, and forced us to consider the limitations of the model.

In Chapter 6 we were able to verify the results of our model of the cultural evolution of language by a direct comparison of the qualitative results obtained against observations of human linguistic diversity. In the following chapter, we took this a step further by conducting a further experiment to verify our findings, using a different model that had already been extensively documented.

That we are able to reproduce our some of our key results in a second, quite different, model gives us significantly more confidence in our findings. These findings are now summarised in the next two sections.

8.3 Redundancy in Language Evolution

Perhaps the first notable finding of this work, in Chapters 4 and 5, was that a redundant language capacity (i.e. one capable of using more signals, or representing more meanings than strictly required) improved the fitness of agents possessing it.

The benefit of a redundant language capacity could also overcome moderate fitness penalties awarded according to capacity.

This was interesting because the benefit occurred despite the absence of noise – the most commonly cited benefit of linguistic redundancy. We concluded that, in this case, the benefit was in allowing greater flexibility for agents attempting to learn conflicting lexicons. It was noted that along a language continuum, agents were able to understand neighbouring agents with a high degree of success despite large differences in dialect along the continuum. While a similar idea has been previously presented (Lass (1997) pointed to redundancy as allowing changes to occur in language without affecting the success of communication), this work has brought into focus the likelihood that redundancy itself provides a significant benefit to language learners, improving their ability to learn language from conflicting and contradictory evidence provided by different speakers.

However, the role and importance of redundancy requires further investigation, and is left as work for the future.

8.4 Language Change and Neutral Evolution

What is probably the most significant contribution of this thesis however, has been the investigation of the evolution of linguistic diversity, and subsequent discussion. This is a topic of current and ongoing debate. One view holds that language diversity and/or change must be adaptive, or otherwise it would not occur for a variety of reasons. Countering this is the view that change and diversity are unavoidable, and entirely to be expected due to the way language is learned. In this debate, our experiments and debate supported this latter case. With regard to simulation work, we argued that improbable assumptions coded into models supporting the former position gave rise to their results. To counter those results, we presented results from two distinct models supporting the latter argument.

Previous work on language change as a form of neutral evolution was presented by Lass (1997) and Nettle (1999a), who presented arguments for and against, respectively. We further added to this existing work, and briefly outlined further arguments forwarding language change as being adaptively neutral. In doing this we also took time to look at possible challenges introduced by related ongoing debates in linguistics, on linguistic relativity and non-uniformitarianism.

As yet, this debate is not thoroughly resolved, and the work presented can be greatly extended. For example, linguistic ecology (Section 2.5.5) would appear to be an area where computational models could be productively applied – investigating more closely the social structures supportive of language stability, of creole genesis, or takeover by one of a number of competing languages.

8.5 Future Directions

As noted at the beginning, this thesis is the product of interdisciplinary research. Such work is often troubled, as has been noted by Bullock (1997) in the conclusions of his PhD thesis. In looking at the evolution of languages, however, some of the greatest conceptual and intellectual struggles have not been in attempting to bridge the gap between computer and linguistic sciences, but between different branches of linguistics itself – where principles and ideas in fields such as historical linguistics and sociolinguistics seem at times to be in opposition. It is to be hoped that, over time, work on the evolution of language might help to draw some of these distinct approaches together, providing a bigger picture that frames the sometimes opposing views that currently exist in, say, socio-linguistics and historical linguistics.

Regarding the work of this thesis, there are a variety of ways in which future research might expand upon it. As noted above, both the role of redundancy in language evolution and the effect of different linguistic ecologies present further questions.

Other interesting, and related, directions for future research include examining how the evolution of languages differs for different aspects of language (e.g. grammar versus phonology), or for examining the role played by different mechanisms of language change. For example Labov (1994, Chapter 18) examines the role played by two different mechanisms of language change, and proposes that both have some part to play. ALife work may be able to explore the workings of such mechanisms more closely, allowing insight not possible from currently available linguistic evidence.

Appendix A – Table of Signals Learned

The table below lists the full set of signals learned by a population in one of the experiments detailed in Section 5.4. This shows that considerable diversity exists, with neighbours tending to use the same – or similar signals. This demonstrates the spontaneous emergence of dialects.









Agent No.	Lang. Nodes	Signal 1	Signal 2	Signal 3	Signal 4	Signal 5	Signal 6	Signal 7	Signal 8
0	6	++++-+	+-----	-+++++	-----	+--+-	-----	++++--	+--+-
1	6	-+---+	+-----	-+++++	-----	++++--	-----	+-----	++++--
2	6	-+---+	+-----	-+++++	-----	++++--	-----	+-----	-+---+
3	6	---++	+-----	---++	-----	++++--	-----	+--+-	-+---+
4	5	---+-	+---+	---+-	-----	++++--	-----	+--+-	-+---+
5	7	---++++	+---+-	---++++	-----	++++--	-----	+-----	-+---+
6	6	---++	+---+-	-----	-----	++++--	+---+	-----	-+---+
7	6	---++	+---+-	-----	+---+	++++--	+---+	-----	-+---+
8	5	---+-	+---+	-+---	+---+	++++--	+---+	---+-	-+---
9	6	---+-	+---+-	-+---	++++--	++++--	+---+	---+-	-+---
10	6	-+---+	---++	-+---+	++++--	++++--	+---+	---++	-+---
11	5	---++	---+-	-+---	++++--	++++--	+---+	+---+	-+---
12	6	+---+-	---++	-+---	++++--	++++--	-----	---++	-+---+
13	6	---++	---+-	---+-	++++--	++++--	+---+	---++	-+---+
14	5	-+---+	---+-	---+-	+---+	++++--	+---+	---+-	+---+
15	6	-+---+	-----	+---+-	+---+-	+---+-	+---+-	---+-	+---+-
16	6	-+---+	+-----	+---+-	---++	+---+-	+---+-	---++	+---+-
17	6	---++	+-----	+---+-	---++	+---+-	+---+-	---++	+---+-
18	6	---++	+-----	---+-	---++	+---+-	-----	---++	---++
19	5	---+-	+---+	---+-	+---+	+---+	+---+	---+-	+---+
20	6	---+-	+---+-	+---+-	---++	+---+-	+---+-	-+---	+---+-
21	5	---+-	+---+	+---+	---++	-----	+---+	-+---	+---+
22	7	---++++	++++--	++++--	+---+-	-----	++++--	+---+-	---++
23	6	-+---+	+---+-	+---+-	+---+-	---+-	++++--	+---+-	+---+-
24	6	-+---+	+---+-	+---+-	+---+-	---+-	++++--	+---+-	---++
25	5	-+---+	+---+-	+---+-	-+---	---+-	++++--	+---+-	---+
26	6	-+---+	+---+-	+---+-	+---+-	-----	++++--	+---+-	---++
27	6	++++--	+---+-	++++--	+---+-	-----	-+---	---+-	+---+-
28	7	+---+-	+---+-	++++--	+---+-	---++	+---+-	---+-	+---+-
29	5	+---+	++++--	++++--	+---+	-+---	+---+	---+-	+---+
30	6	+---+-	++++--	++++--	+---+-	-+---	+---+-	-+---	+---+-
31	6	---+-	++++--	+---+-	+---+-	-+---	+---+-	-+---	+---+-
32	6	-----	++++--	-+---	+---+-	-+---	---+-	---+-	+---+-
33	7	-+---+	+---+-	+---+-	+---+-	-+---	---++	---+-	+---+-
34	6	-+---+	+---+-	+---+-	+---+-	-+---	---++	---+-	+---+-
35	5	-+---	+---+	+---+	+---+	-+---	-----	+---+	+---+
36	6	-+---+	+-----	+---+-	-+---	-+---	-----	+---+-	+---+-
37	6	-+---+	+-----	+---+-	-+---	-----	-----	+---+-	+---+-
38	6	+---+-	-----	+---+-	-+---	---+-	+---+-	+---+-	+---+-
39	7	+---+-	-+---	-+---	-+---	---++	+---+-	+---+-	+---+-
40	6	---+-	-+---	-+---	-----	---++	+---+-	+---+-	+---+-
41	6	---+-	-+---	-----	-----	---++	+---+-	+---+-	+---+-
42	7	---+-	-+---	-+---	-----	---++	+---+-	---++	+---+-

Agent No.	Lang. Nodes	Signal 1	Signal 2	Signal 3	Signal 4	Signal 5	Signal 6	Signal 7	Signal 8
97	6	++++--	++----	++++++	+-----	+--+--	-++++-	-++---	+-----
98	5	++++-	-----	+++++	+-----	+--++	-++++	-++---	+-----
99	7	+++++-	+-----	-+++++	+-----	+--+--	-++++-	-++---	+-----

Table A.1: This table shows the full set of the signals from one set of results from one of experiments shown in section 5.4

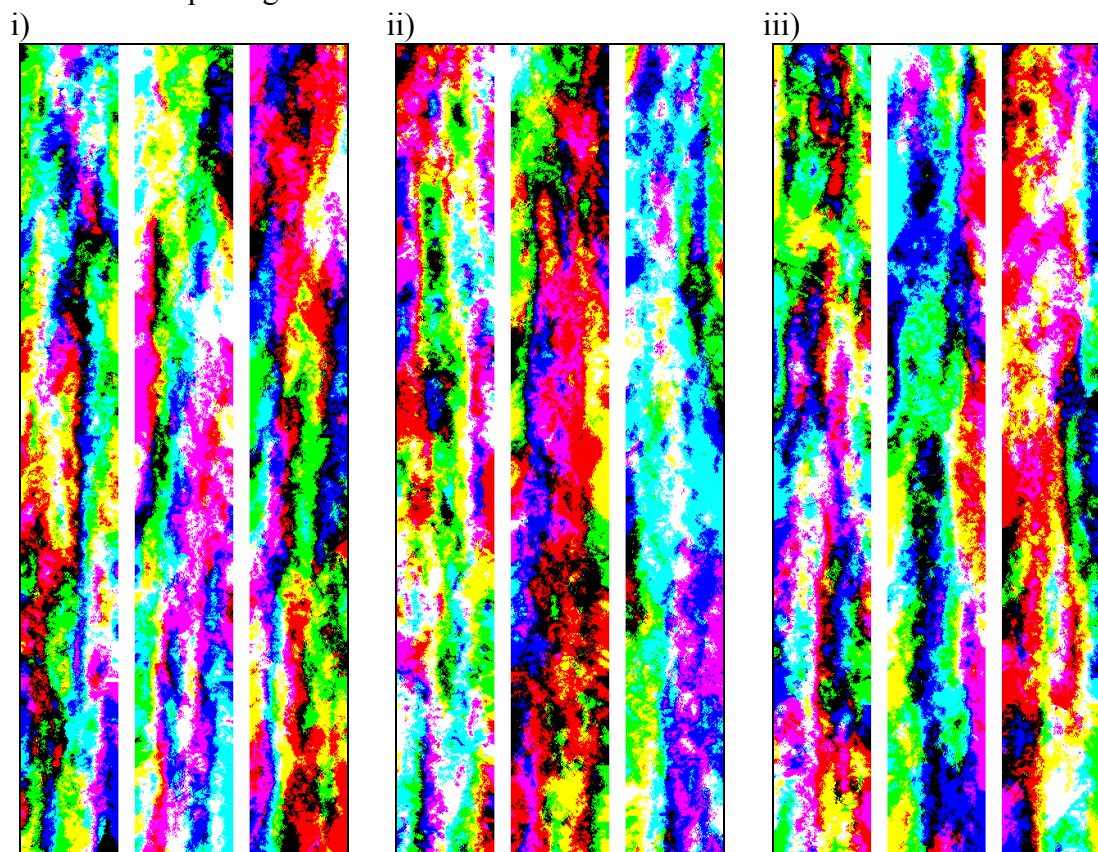
Appendix B – Colour Plates

Colour Plate 6.1 Signals and colour values


Signal	Colour	Colour Name	Signal	Colour	Colour Name
0 0 0		Black	1 0 0		Red
0 0 1		Blue	1 0 1		Magenta
0 1 0		Green	1 1 0		Yellow
0 1 1		Turquoise	1 1 1		White

Colour Plate 6.2 Evolution of Communication Schemes


(a).i-ix Set of results after agents learn only from agents in the parent generation. Shown in each are the last 1000 generations of 100,000. See section 6.6.2 for more detail on interpreting the results.




iv)




v)




vi)




vii)




viii)

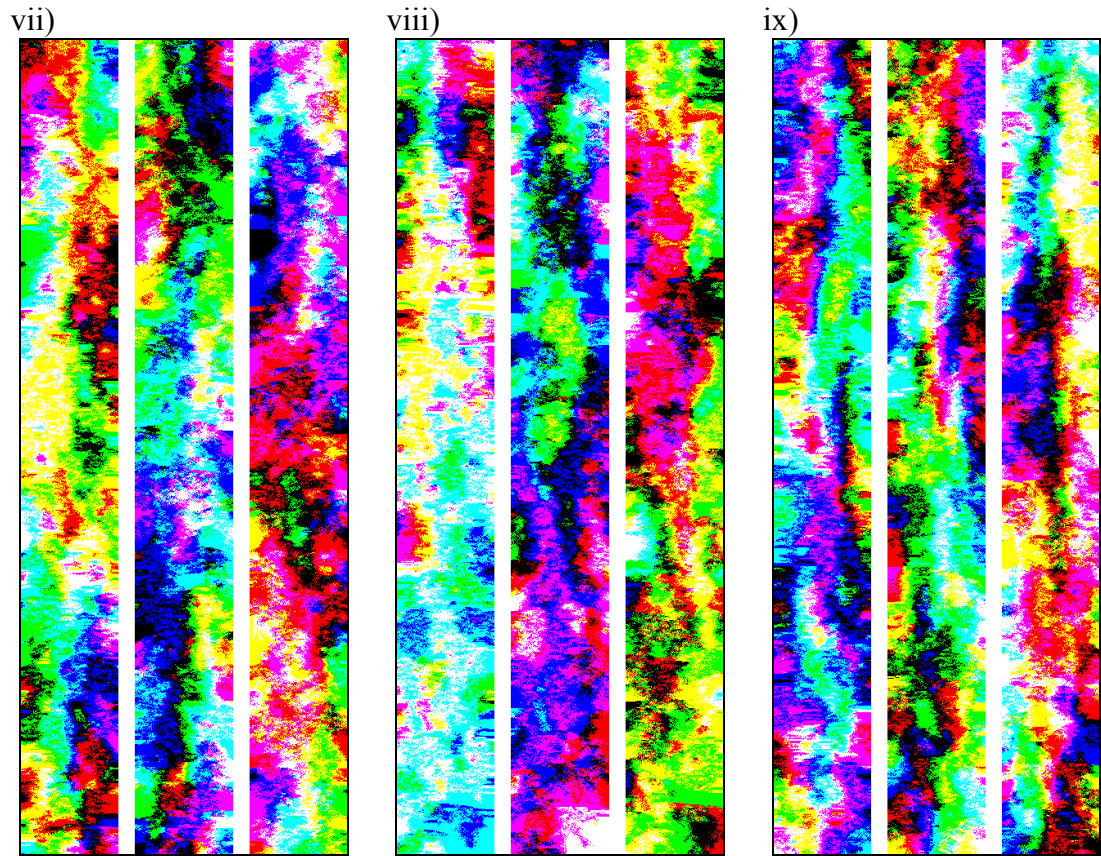


ix)




(b).i-ix Set of results after agents learn from other agents in their own generation as well as from agents in the parent generation. Shown in each are the last 1000 generations of 100,000.






Colour Plate 6.3 Emergence of global dialects

i) Convergence ii) and iii) Global dialects, with internal variation.




Colour Plate 6.4 Diversity from Homogeneity


(a) From homogenous signalling schemes to diversity, three examples of the effect of noisy signals at 0.1% noise (i-iii) and three with 0.01% noise (iv-vi). Shown in each are the first 1000 generations of a run.



(b) Six examples of the effect of bias from initial weights on signal evolution. Initial weight values range from -0.05 to 0.05 (i-iii) or from -0.005 to 0.005 (iv-vi). Shown in each are the first 1000 generations of a run.



(c) A further example of the effect of innate bias. i – From homogeneity to diversity, the first thousand generations. ii – No evidence of original homogeneity, from the 9000th to the 10,000th generation.



References

- Adami, C., R. K. Belew, H. Kitano and C. E. Taylor (1998). Artificial Life VI. In Artificial Life VI, UCLA, MIT Press.
- Adamska-Salaciak, A. (1997). Sound Laws: Reactions Past and Present. In Historical Linguistics 1997. M. S. Schmid, J. R. Austin and D. Stein (Eds.). Amsterdam, John Benjamins: 1-13.
- Aiello, L. C. (1996). Terrestriality, Bipedalism, and the Origin of Language. In Evolution of Social Behaviour Patterns in Primates and Man (Proceedings of the British Academy 88). E. G. Runciman, J. Maynard-Smith and D. R.I.M. (Eds.). Oxford, Oxford University Press: 269–289.
- Aiello, L. C. (2002). Fossil and Other Evidence for the Origin of Language. In J. Hurford and T. Fitch (Eds.), proceedings of the Fourth International Conference on the Evolution of Language, Harvard University, Cambridge, M.A., U.S.A.: 2.
- Aitchison, J. (1991). Language Change: Progress or Decay?, Cambridge University Press.
- Aitchison, J. (1998). On Discontinuing the Continuity-Discontinuity Debate. In Approaches to the Evolution of Language. J. R. Hurford, M. Studdert-Kennedy and C. Knight (Eds.), Cambridge University Press: 17-29.
- Anderson, H. (1973). Abductive and Deductive Change. Language **49**: 765-93.
- Arita, T. and Y. Koyama (1998). Evolution of Linguistic Diversity in a Simple Communication System. In C. Adami, R. K. Belew, H. Kitano and C. E. Taylor (Eds.), Artificial Life IV, UCLA, MIT Press: 9-17.
- Arita, T. and C. E. Taylor (1996). A Simple Model for the Evolution of Communication. In L. J. Fogel, P. J. Angeline and T. Bäck (Eds.), Evolutionary Programming V, MIT Press: 405-409.
- Armstrong, D. F., W. C. Stokoe and S. E. Wilcox (1994). Signs of the Origin of Syntax. Current Anthropology **35**(4): 349-368.
- Axelrod, R. (1984). The Evolution of Cooperation. New York, Basic Books.
- Axelrod, R. (1997a). The Complexity of Cooperation: Agent Based Models of Competition and Collaboration. Princeton, New Jersey, Princeton University press.
- Axelrod, R. (1997b). The Dissemination of Culture: A Model with Local Convergence and Global Polarization. Journal of Conflict Resolution **41** (1997): 203-226.
- Barlow, H. B. (2001). Redundancy Reduction Revisited. Network: Computation in Neural Systems **12**: 241-253.

- Batali, J. (1994). Innate Biases and Critical Periods: Combining Evolution and Learning in the Acquisition of Syntax. In R. Brooks and P. Maes (Eds.), Proceedings of the Fourth Artificial Life Workshop, Cambridge, MA, MIT Press: 160-171.
- Batali, J. (1998). Computational Simulations of the Emergence of Grammar. In Approaches to the Evolution of Language. J. Hurford, M. Studdert-Kennedy and C. Knight (Eds.), Cambridge University Press: 405-426.
- Bates, E. and J. C. Goodman (1999). On the Emergence of Grammar from the Lexicon. In The Emergence of Language. B. MacWhinney (Ed.). Mahwah, New Jersey, Lawrence Erlbaum Associates: 29-80.
- Baugh, A. C. and T. Cable (1978). A History of the English Language. London, Routledge.
- Belew, R. K., J. McInerney and N. N. Schraudolph (1991). Evolving Networks: Using the Genetic Algorithm with Connectionist Learning. In Artificial Life II. C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), Addison-Wesley.
- Bennet, S., S. McRobb and R. Farmer (2001). Object-Oriented Systems Analysis and Design Using Uml, McGraw-Hill.
- Bichakjian, B. H. (1999). Language Evolution and the Complexity Criterion. Psychology **10**(33).
- Bickerton, D. (1981). Roots of Language. Ann Arbor, Michigan, Karoma.
- Bickerton, D. (1984). The Language Bioprogram Hypothesis. Behavioral and Brain Sciences **7**: 173-221.
- Bickerton, D. (1990). Species and Language. Chicago, Chicago University Press.
- Brighton, H. and S. Kirby (2001). The Survival of the Smallest: Stability Conditions for the Cultural Evolution of Compositional Language. In J. Kelmen and P. Susik (Eds.), 5th European Conference on Artificial Life, Prague, Published in Advances in Artificial Life, Springer, Berlin: 592-601.
- Briscoe, E. J. (2000a). Evolutionary Perspectives on Diachronic Syntax. In Diachronic Syntax: Models and Mechanisms. S. Pintzuk, G. Tsoulas and A. Warner (Eds.). Oxford, Oxford University Press: 75-108.
- Briscoe, E. J. (2000b). Grammatical Acquisition: Inductive Bias and Coevolution of Language and the Language Acquisition Device. Language **76**(2): 245-296.
- Briscoe, E. J. (2000c). Macro and Micro Models of Linguistic Evolution. In J.-L. Dessalles and L. Ghadapkpour (Eds.), The Evolution of Language, 3rd International Conference, Paris: 27-29.
- Buckley, C. and J. Steele (2002). Evolutionary Ecology of Spoken Language: Co-Evolutionary Hypotheses Are Testable. World Archaeology **34**(1): 26-46.

- Bullock, S. (1997). Evolutionary Simulation Models: On Their Character, and Application to Problems Concerning the Evolution of Natural Signalling Systems, PhD Thesis. University of Sussex, Brighton
- Bullock, S. (2000). Something to Talk About: Conflict and Coincidence of Interest in the Evolution of Shared Meaning. In J.-L. Dessalles and L. Ghadapkpour (Eds.), The Evolution of Language, 3rd International Conference, Paris.
- Burling, R. (1998). Comprehension and Production in Early Language. In The evolution of language, 2nd international conference, London.
- Burling, R. (2000). The Slow Growth of Language in Children. In J.-L. Dessalles and L. Ghadapkpour (Eds.), The Evolution of Language, 3rd international conference, Paris: 35.
- Burnett, J. (1787). Of the Origin and Progress of Language.
- Calvin, W. H. and D. Bickerton (2000). Lingua Ex Machina: Reconciling Darwin and Chomsky with the Human Brain. Cambridge, MA, MIT Press.
- Cangelosi, A., A. Greco and S. Harnad (2002). Symbol Grounding and the Symbolic Theft Hypothesis. In Simulating the Evolution of Language. A. Cangelosi and D. Parisi (Eds.). London, Springer: 191-210.
- Cangelosi, A. and S. Harnad (2000). The Adaptive Advantage of Symbolic Theft over Sensorimotor Toil: Grounding Language in Perceptual Categories. Evolution of Communication 4(1): 117-142.
- Cangelosi, A. and D. Parisi (1998). The Emergence of a "Language" in an Evolving Population of Neural Networks. Connection Science 10(2): 83-97.
- Cangelosi, A. and D. Parisi, Eds. (2002). Simulating the Evolution of Language. London, Springer.
- Carstairs-McCarthy, A. (1999). The Origins of Complex Language: An Inquiry into the Evolutionary Beginnings of Sentences, Syllables, and Truth. Oxford, Oxford University Press.
- Caruna, R. and J. D. Schaffer (1988). Representation and Hidden Bias: Gray Vs. Binary Coding for Genetic Algorithms. In Fifth International Conference on Machine Learning, Morgan Kaufmann.
- Cavalli-Sforza, L. L. and M. W. Feldman (1978). Towards a Theory of Cultural Evolution. Interdisciplinary Review 3: 99-107.
- Cavalli-Sforza, L. L. and M. W. Feldman (1981). Cultural Transmission and Change: A Quantitative Approach, Princeton University Press.
- Chambers, J. K. (1995). Sociolinguistic Theory, Blackwell.
- Chambers, J. K. and P. Trudgill (1980). Dialectology, Cambridge University Press.

- Chambers, J. K. and P. Trudgill (1998). Dialectology, 2nd Edition, Cambridge University Press.
- Chomsky, N. (1972). Language and Mind. San Diego, Harcourt Brace Jovanovich.
- Christiansen, M. H. (2000). Using Artificial Language Learning to Study Language Evolution: Exploring the Emergence of Word Order Universals. In J.-L. Dessalles and L. Ghadapkpour (Eds.), The Evolution of Language, 3rd International Conference, Paris: 45-48.
- Clark, D. A., P. P. Mitra and S. S.-H. Wang (2001). Scalable Architecture in Mammalian Brains. Nature **411**: 183-193.
- Clark, R. A. J. (1996). Internal and External Factors Affecting Language Change: A Computational Model, MSc Dissertation. University of Edinburgh
- Crystal, D. (1987). The Cambridge Encyclopedia of Language, Cambridge University Press.
- Dautenhahn, K. and C. Nehaniv, Eds. (1999). Proceedings of the Aisb '99 Symposium on Creative Evolutionary Systems. Edinburgh, The Society for the Study of Artificial Intelligence and Simulation of Behaviour.
- Dawkins, R. (1976). The Selfish Gene.
- de Boer, B. (1997). Generating Vowel Systems in a Population of Agents. In 4th European Conference on Artificial Life, ECAL 97, Brighton, UK.
- de Boer, B. (2000). Emergence of Sound Systems through Self-Organisation. In The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form. C. Knight, M. Studdert-Kennedy and J. R. Hurford (Eds.), Cambridge University Press: 177-198.
- de Boer, B. and P. Vogt (1999). Emergence of Speech Sounds in a Changing Population. In D. Floreano, J.-D. Nicoud and F. Mondada (Eds.), Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life, ECAL 99, Lausanne, Switzerland, Springer: 664-673.
- Deacon, T. (1997). The Symbolic Species: The Co-Evolution of Language and the Human Brain, W.H. Norton & Co., USA and Allen Lane, The Penguin Press, UK.
- Deacon, T. W. (1992). Brain-Language Coevolution. In The Evolution of Human Languages. J. A. Hawkins and M. Gell-Mann (Eds.), Addison-Wesley: 273-303.
- Dessalles, J.-L. and L. Ghadapkpour, Eds. (2000). Proceeding of the Evolution of Language, 3rd International Conference. Paris.
- Di Paolo, E. (1996). Some False Starts in the Construction of a Research Methodology for Artificial Life. In J. Noble and S. Parsowith (Eds.), The 9th White House Papers: Graduate Research in the Cognitive and Computing

Sciences at Sussex, University of Sussex, School of Cognitive and Computing Sciences.

- Di Paolo, E. (1997a). An Investigation into the Evolution of Communication. Adaptive Behavior **6**(2): 285-324.
- Di Paolo, E. (1997b). Social Coordination and Spatial Organization: Steps Towards the Evolution of Communication. In ECAL 97, Brighton, UK.
- Di Paolo, E. (1999a). A Little More Than Kind and Less Than Kin: The Unwarrented Use of Kin Selection in Spatial Models of Communication. In D. Floreano, J.-D. Nicoud and F. Mondada (Eds.), *Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life, ECAL 99, Lausanne, Switzerland, Springer: 504-513.*
- Di Paolo, E. (1999b). *On the Evolutionary and Behavioral Dynamics of Social Coordination: Models and Theoretical Aspects*. University of Sussex
- Dircks, C. and S. C. Stoness (1999). Investigating Language Change: A Multi-Agent Neural-Network Based Simulation. In D. Floreano, J.-D. Nicoud and F. Mondada (Eds.), *Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life, ECAL 99, Lausanne, Switzerland, Springer: 720-724.*
- Doran, J. (1996). *Simulating Societies Using Distributed Artificial Intelligence*. In Social Science Microsimulation. K. G. Troitzsch, U. Mueller, G. N. Gilbert and J. E. Doran (Eds.). Berlin, Springer: 381-393.
- Dunbar, R. (1996). Grooming, Gossip and the Evolution of Language, Faber & Faber.
- Elliot, A. J. (1981). Child Language. Cambridge, U.K., Cambridge University Press.
- Elman, J. L. (1993). Learning and Development in Neural Networks: The Importance of Starting Small. Cognition(48): 71-99.
- Elman, J. L. (1999). The Emergence of Language: A Conspiracy Theory. In The Emergence of Language. B. MacWhinney (Ed.). Mahwah, New Jersey, Lawrence Erlbaum Associates: 1-28.
- Epstein, J. M. and R. Axtell (1996). Growing Artificial Societies: Social Science from the Bottom Up, Brookings Institution Press, Washington, D.C., The MIT Press, Cambridge, MA.
- Forrester, J. W. (1971). World Dynamics. Cambridge, MA, MIT Press.
- Frankfort-Nachmias, C. and D. Nachmias (1992). Research Methods in the Social Sciences. London, Edward Arnold.
- Fyfe, C. and D. Livingstone (1997). Developing a Community Language. In 4th European Conference on Artificial Life, ECAL 97, Brighton, UK.

- Gibson, E. and K. Wexler (1994). Triggers. Linguistic Inquiry(25).
- Gilbert, N. and K. G. Troitzsch (1999). Simulation for the Social Scientist, Open University Press.
- Gleick, J. (1987). Chaos, Heinmann.
- Gopnik, M. and M. B. Crago (1991). Familial Aggregation of a Developmental Language Disorder. Cognition **39**: 1-50.
- Gould, S. J. and E. S. Vrba (1982). Exaptation - a Missing Term in the Science of Form. Paleobiology **8**(1): 4-15.
- Grafen, A. (1991). Modelling in Behavioural Ecology. In Behavioural Ecology. J. R. Krebs and N. B. Davies (Eds.). Oxford, Blackwell Scientific Publications: 5-31.
- Gray, F. (1953). Pulse Code Communication. U.S. Patent Number 2,632,058.
- Grice, H. P. (1969). Utterer's Meaning and Intention. Philosophical Review **68**: 147-177.
- Grice, H. P. (1975). Logic and Conversation. In Syntax and Semantics Volume 3: Speech Acts. P. Cole and J. Morgan (Eds.). NY, Academic Press.
- Grimm, J. (1822). Deutsche Grammatik, Part I: Zweite Ausgabe. Göttingen, Dieterich.
- Hamilton, W. D. (1964). The Genetical Evolution of Social Behavior. Journal of Theoretical Biology **7**: 1-52.
- Harnad, S. (1990). The Symbol Grounding Problem. Physica D **42**: 335-346.
- Haugen, E. (1971). The Ecology of Language. In Ecology of Language, Essays by Einar Haugen. A. S. Dil (Ed.). Stanford, CA, Stanford University Press: 325-339.
- Hauser, M. D. (1996). The Evolution of Communication. Cambridge, Massachusetts, MIT Press.
- Hauser, M. D. and P. Marler (1999). Animal Communication. In The Mit Encyclopedia of the Cognitive Sciences. R. A. Wilson and F. C. Keil (Eds.). Cambridge, MA, MIT Press: 22-23.
- Haykin, S. (1998). Neural Networks. New Jersey, Prentice Hall.
- Hinton, G. E. and Z. Ghahramani (1997). Generative Models for Discovering Sparse Distributed Representations. Philosophical transactions fo the Royal Society **B**(352): 1177-1190.
- Hockett, C. F. (1960). The Origin of Speech. Scientific American(203): 88-93.

- Hofbauer, J. and K. Sigmund (1988). The Theory and Evolution of Dynamical Systems. Cambridge, Cambridge University Press.
- Holland, J. H. (1975). Adaptation in Natural and Artificial Systems. Ann Arbor, MI, University of Michigan Press (2nd ed.: MIT Press, 1992).
- Holland, J. H. (1998). Emergence; from Chaos to Order. Oxford, Oxford University Press.
- Hurford, J. R. (1992). An Approach to the Phylogeny of the Language Faculty. In The Evolution of Human Languages. J. A. Hawkins and M. Gell-Mann (Eds.), Addison-Wesley: 273-303.
- Hurford, J. R. (1996). Functional Innateness: Explaining the Critical Period for Language Acquisition. In Submitted to Functionalism/Formalism in Linguistics (Proceedings of the 23rd Annual Uwm Linguistics Symposium).
- Hurford, J. R. (1999). The Evolution of Language and Languages. In The Evolution of Culture. R. Dunbar, C. Knight and C. Power (Eds.), Edinburgh University Press.
- Hurford, J. R., M. Studdert-Kennedy and C. Knight (1998). Approaches to the Evolution of Language, Cambridge University Press.
- Hutchins, E. and B. Hazelhurst (1995). How to Invent a Lexicon: The Development of Shared Symbols in Interaction. In Artificial Societies: The Computer Simulation of Social Life. N. Gilbert and R. Conte (Eds.), UCL Press.
- Huynen, M. A. (1995). Exploring Phenotype Space through Neutral Evolution, Technical Report Preprint, Santa Fe Institute
- Johnson, K. (1997). Acoustic and Auditory Phonetics. Cambridge, MA., Blackwell.
- Kegl, J. and G. A. Iwata (1989). Language De Signos Nicaraguense: A Pidgin Sheds Light on the "Creole?" Asl. In proceedings of the fourth annual meeting of the Pacific Linguistics conference, University of Oregon, Eugene, Ore., U.S.A.
- Kimura, M. (1983). The Neutral Theory of Molecular Evolution, Cambridge University Press.
- King, R. D. (1969). Push Chains and Drag Chains. Glossa 3(1): 3-21.
- Kirby, S. (1998). Fitness and the Selective Adaptation of Language. In Approaches to the Evolution of Language. J. R. Hurford, M. Studdert-Kennedy and C. Knight (Eds.), Cambridge University Press: 359-383.
- Kirby, S. (1999). Function, Selection, and Innateness: The Emergence of Language Universals. Oxford, Oxford University Press.
- Kirby, S. (2000). Syntax without Natural Selection: How Compositionality Emerges from Vocabulary in a Population of Learners. In The Evolutionary Emergence

- of Language: Social Function and the Origins of Linguistic Form. C. Knight, M. Studdert-Kennedy and J. R. Hurford (Eds.), Cambridge University Press.
- Kirchkamp, O. (1996). Spatial Evolution of Automata in the Prisoner's Dilemma. In Social Science Microsimulation. K. G. Troitzsch, U. Mueller, G. N. Gilbert and J. E. Doran (Eds.). Berlin, Springer: 307-358.
- Kitano, H. (1990). Designing Neural Networks Using Genetic Algorithms with Graph Generation System. Complex Systems 4: 461-476.
- Knuth, D. (1969). The Art of Computer Programming, Volume 1: Fundamental Algorithms. Reading, Ma., Addison-Wesley.
- Krebs, J. R. and R. Dawkins (1984). Animal Signals: Mind Reading and Manipulation. In Behavioural Ecology. J. R. Krebs and N. B. Davies (Eds.), Blackwell Scientific Publications.
- Kroch, A. (1989). Reflexes of Grammar in Patterns of Language Change. Journal of Language Variation and Change 1: 199-244.
- Labov, W. (1972). Sociolinguistic Patterns. Philadelphia, University of Pennsylvania Press.
- Labov, W. (1994). Principles of Linguistic Change: Internal Factors. Oxford, Blackwell.
- Langton, C. G. (1989). Preface in Artificial Life. In S.F.I. Studies in the Sciences of Complexity. C. G. Langton (Ed.), Addison-Wesley.
- Langton, C. G. (1991). Preface. In Artificial Life II. C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.). Redwood City, CA, Addison-Wesley: xiii-xviii.
- Lass, R. (1997). Historical Linguistics and Language Change. Cambridge, UK, Cambridge University Press.
- Latané, B. (1981). The Psychology of Social Impact. American Psychologist 36: 343-365.
- Lee, M. G. (2000). The Ethics of Deception: Why Ai Must Study Selfish Behaviour. AISB Quaterly(104): 20-27.
- Lieberman, P. (1992). On the Evolution of Human Language. In The Evolution of Human Languages. J. A. Hawkins and M. Gell-Mann (Eds.), Addison-Wesley: 21-47.
- Lightfoot, D. (1999). The Development of Language: Acquisition, Change and Evolution. Malden, Massachusetts, Blackwell Publishers Ltd.
- Lindblom, B. (1998). Systemic Constraints and Adaptive Change in the Formation of Sound Structure. In Approaches to the Evolution of Language. J. R. Hurford,

- M. Studdert-Kennedy and C. Knight (Eds.), Cambridge University Press: 242-264.
- Livingstone, D. (1999a). Methodology in Madness? Computing and Information Systems, CIS Departmental Journal, University of Paisley 6(3): 113-118.
- Livingstone, D. (1999b). On Modelling the Evolution of Language and Languages. In GECCO-99 Student Workshop, Orlando.
- Livingstone, D. (2000a). Computational Models of Language Change and Diversity. In B. deBoer and P. Vogt (Eds.), presented at Bene-Evolang 2000, the Belgian-Dutch conference on the evolution of language, Brussels.
- Livingstone, D. (2000b). A Modified-Neutral Theory for the Evolution of Linguistic Diversity. In The 3rd International Conference on the Evolution of Language, Paris.
- Livingstone, D. (2002). The Evolution of Dialect Diversity. In Simulating the Evolution of Language. A. Cangelosi and D. Parisi (Eds.). London, Springer-Verlag.
- Livingstone, D. and C. Fyfe (1998a). A Computational Model of Language-Physiology Coevolution. Computing and Information Systems, CIS Departmental Journal, University of Paisley 5(2): 55-62.
- Livingstone, D. and C. Fyfe (1998b). A Computational Model of Language-Physiology Coevolution. In The 2nd International Conference on the Evolution of Language, London.
- Livingstone, D. and C. Fyfe (1999a). Dialect in Learned Communication. In K. Dautenhahn and C. Nehaniv (Eds.), AISB '99 Symposium on Imitation in Animals and Artifacts, Edinburgh, The Society for the Study of Artificial Intelligence and Simulation of Behaviour: 139-146.
- Livingstone, D. and C. Fyfe (1999b). Modelling the Evolution of Linguistic Diversity. In D. Floreano, J.-D. Nicoud and F. Mondada (Eds.), Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life, ECAL 99, Lausanne, Switzerland, Springer: 139-146.
- Livingstone, D. and C. Fyfe (2000). Modelling Language-Physiology Coevolution. In The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form. C. Knight, M. Studdert-Kennedy and J. R. Hurford (Eds.), Cambridge University Press: 199-215.
- Lorenz, K. (1966). Evolution of Ritualization in the Biological and Cultural Spheres. Philosophical Transactions of the Royal Society of Britain 251: 273-284.
- Lucy, J. A. (1999). Linguistic Relativity Hypothesis. In The Mit Encyclopedia of the Cognitive Sciences. R. A. Wilson and F. C. Keil (Eds.). Cambridge, M.A., MIT Press: 475-476.

- MacLennan, B. (1991). Synthetic Ethology: An Approach to the Study of Communication. In Artificial Life II. C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), Addison-Wesley.
- MacLennan, B. J. and G. M. Burghardt (1994). Synthetic Ethology and the Evolution of Cooperative Communication. Adaptive Behavior 2: 161-187.
- MacWhinney, B., Ed. (1999). The Emergence of Language, Lawrence Erlbaum Associates.
- Maeda, Y., T. Sasaki and M. Tokoro (1997). Self-Reorganizing of Language Triggered by "Language Contact". In ECAL 97, Brighton, UK.
- Maess, B., S. Koelsch, T. C. Gunter and A. D. Friederici (2001). Musical Syntax Is Processed in Broca's Area: An Meg Study. Nature Neuroscience 4: 540-545.
- Marenbon, J. (1987). English Our English: The New Orthodoxy Examined. London, Centre for Policy Studies. Extract reprinted in Proper English? T. Crowley. London, Routledge, 1991
- Miller, G. F. (1995). Artificial Life as Theoretical Biology: How to Do Real Science with Computer Simulation, Cognitive Science Research Paper. Brighton, Cognitive and Computing Sciences, University of Sussex
- Miller, G. F., P. M. Todd and S. U. Hegde (1989). Designing Neural Networks Using Genetic Algorithms. In D. Schaffer (Ed.), Proceedings of the Third International Conference on Genetic Algorithms, Morgan Kaufman.
- Milroy, J. (1992). Linguistic Variation and Change. Oxford, Blackwell.
- Milroy, J. (1993). On the Social Origins of Language Change. In Historical Linguistics: Problems and Perspectives. C. Jones (Ed.). London, New York, Longman.
- Milroy, J. and L. Milroy (1993). Mechanisms of Change in Urban Dialects: The Role of Class, Social Network and Gender. International Journal of Applied Linguistics 3(1).
- Milroy, L. (1980). Language and Social Networks. Oxford, Basil Blackwell.
- Mitchell, M. (1996). An Introduction to Genetic Algorithms. Cambridge, MA., MIT Press.
- Mitchell, M. (1999). Evolutionary Computation. In The Mit Encyclopedia of the Cognitive Sciences. R. A. Wilson and F. C. Keil (Eds.). Cambridge, MA, MIT Press: 293-295.
- Mithen, S. (1996). The Prehistory of the Mind. London, Thames and Hudson Ltd.
- Mufwene, S. S. (1997). What Research on Creole Genesis Can Contribute to Historical Linguistics. In Historical Linguistics 1997. M. S. Schmid, J. R. Austin and D. Stein (Eds.). Amsterdam, John Benjamins.

- Mühlenbein, H. and D. Schlierkamp-Voosen (1993). Predictive Models for the Breeder Genetic Algorithm. Evolutionary Computation **1**(1): 25-49.
- Mühlhäusler, P. (1996). Linguistic Ecology: Language Change and Linguistic Imperialism in the Pacific Region. London, Routledge.
- Muller, G. (1996). Exploring and Testing Theories: On the Role of Parameter Optimization in Social Science Computer Simulation. In Social Science Microsimulation. K. G. Troitzsch, U. Mueller, G. N. Gilbert and J. E. Doran (Eds.). Berlin, Springer: 66-75.
- Nettle, D. (1999a). Linguistic Diversity, Oxford University Press.
- Nettle, D. (1999b). Using Social Impact Theory to Simulate Language Change. Lingua **108**: 95-117.
- Nettle, D. and R. I. M. Dunbar (1997). Social Markers and the Evolution of Reciprocal Exchange. Current Anthropology **38**(1): 93-99.
- Newmeyer (2000). The Evolution of Language, 3rd International Conference. In J.-L. Dessalles and L. Ghadapkpour (Eds.), The Evolution of Language, 3rd International Conference, Paris.
- Nimwegen, E. v., J. P. Crutchfield and M. Huynen (1999). Neutral Evolution of Mutational Robustness. Santa Fe Institute working paper.
- Niyogi, P. (2002). Models of Cultural Evolution and Their Application to Language Change. In Linguistic Evolution through Language Acquisition: Formal and Computational Models. E. J. Briscoe (Ed.), Cambridge University Press.
- Niyogi, P. and R. C. Berwick (1995). The Logical Problem of Language Change, Massachusetts Institute of Technology
- Noble, J. (1997). The Scientific Status of Artificial Life. In ECAL 97, Brighton, UK.
- Noble, J. (2000). Defining Animal Communication, and Why It Matters for Understanding the Evolution of Language. In J.-L. Dessalles and L. Ghadapkpour (Eds.), The Evolution of Language, 3rd International Conference, Paris.
- Noble, J. and D. Cliff (1996). On Simulating the Evolution of Communication. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack and S. W. Wilson (Eds.), From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour, MIT Press, Cambridge, MA.
- Oliphant, M. (1996). The Dilemma of Saussurean Communication. BioSystems **37**(1-2): 31-38.
- Oliphant, M. (1997). Formal Approaches to Innate and Learned Communication: Laying the Foundation for Language, PhD. University of California, San Diego

- Ormerod, P. (1994). The Death of Economics. London, Faber and Faber.
- Pagel, M. (2000). The History, Rate and Pattern of World Linguistic Evolution. In The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form. C. Knight, M. Studdert-Kennedy and J. R. Hurford (Eds.), Cambridge University Press: 391-416.
- Pinker, S. (1994). The Language Instinct, Penguin.
- Pinker, S. and P. Bloom (1992). Natural Language and Natural Selection. In The Adapted Mind. J. H. Barkow, L. Cosmides and J. Tooby (Eds.), Oxford University Press.
- Quine, W. V. O. (1987). Quiddities: An Intermittently Philosophical Dictionary. Cambridge, Mass., Harvard University Press.
- Renfrew, C. (1998). The Origins of World Linguistic Diversity: An Archaeological Perspective. In The Origin and Diversification of Language. N. G. Jablonski and L. C. Aiello (Eds.), University of California Press: 171-192.
- Reynolds, C. W. (1987). Flocks, Herds and Schools: A Distributed Behavioral Model. Computer Graphics **21**(4).
- Ronald, E. M. A., M. Sipper and M. S. Capcarrère (1999). Design, Observation, Surprise! A Test of Emergence. Artificial Life **5**(3): 225-239.
- Rousseau, J.-J. (1755). Second Discourse (Discourse on the Origins of Inequality). St. Martin's Press Translation, 1964 (in the Collected Writings of Rousseau, Vol. 3, Eds. R.D. Masters and C. Kelly).
- Sanders, I. (2002). Human Tongue, Pharynx and Vocal Fold Muscles Contain Slow Tonic Muscle, a Distinct Class of Muscle That May Have Evolved for Speech. In J. Hurford and T. Fitch (Eds.), proceedings of the Fourth International Conference on the Evolution of Language, Harvard University, Cambridge, M.A., U.S.A.
- Sapir, E. (1949). The Selected Writings of Edward Sapir. In Language, Culture and Personality. D. G. Mandelbaum (Ed.). Berkely and Los Angeles, University of California Press.
- Saussure, F. D. (1916). Cours De Linguistique Générale. Paris.
- Savage-Rumbaugh, E. S. (2000). Plenary Talk. In J.-L. Dessalles and L. Ghadapkpour (Eds.), presented at The Evolution of Language, 3rd International Conference, Paris.
- Schelling, T. C. (1978). Micromotives and Macrobehavior. New York, W.W. Norton & Company.
- Shannon, C. E. and W. Weaver (1949). A Mathematical Theory of Communication. Urbana, Illinois, University of Illinois Press.

- Skoyles, J. R. (2000). The Singing Origin of Speech. In *The Evolution of Language*, 3rd international conference, Paris: 206-209.
- Slobin, D. I. (1979). Psycholinguistics. Gelview, Illinois, Scott, Foresman and Company.
- Slobin, D. I. (2003). Language and Thought Online: Cognitive Consequences of Linguistic Relativity. In Language in Mind: Advances in the Investigation of Language and Thought. D. Gentner and S. Goldin-Meadow (Eds.). Cambridge, MA, Bradford Books, MIT Press.
- Smith, K. (2002). The Cultural Evolution of Communication in a Population of Neural Networks. Connection Science.
- Stearns, S. C. and R. F. Hoekstra (2000). Evolution: An Introduction. Oxford, Oxford University Press.
- Steels, L. (1996a). Emergent Adaptive Lexicons. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack and S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour*, MIT Press, Cambridge, MA.
- Steels, L. (1996b). A Self-Organizing Spatial Vocabulary. Artificial Life 2(4).
- Steels, L. (1996c). Self-Organizing Vocabularies. In *Proceedings of Artificial Life V*, Cambridge, MA, MIT Press: 179-184.
- Steels, L. (1998). Synthesising the Origins of Language and Meaning Using Co-Evolution, Self-Organisation and Level Formation. In Approaches to the Evolution of Language. J. R. Hurford, M. Studdert-Kennedy and C. Knight (Eds.), Cambridge University Press: 384-404.
- Steels, L. and F. Kaplan (1998). Stochasticity as a Source of Innovation in Language Games. In C. Adami, R. K. Belew, H. Kitano and C. E. Taylor (Eds.), *Artificial Life VI*, UCLA, MIT Press: 368-376.
- Steels, L. and P. Vogt (1997). Grounding Adaptive Language Games in Robotic Agents. In *4th European Conference on Artificial Life, ECAL 97*.
- Sternberg (1996). Cognitive Psychology. Fort Worth, :Harcourt Brace.
- Storey, N. (1996). Safety-Critical Computer Systems, Addison-Wesley.
- Swift, J. (1712). Extracts from *a Proposal for Correcting, Improving and Ascertaining the English Language*. In Reprinted in Proper English? T. Crowley (Ed.). London, Routledge, 1991.
- Tannen, D. (1994). Talking from 9 to 5, William Morrow & Company Inc.
- Taylor, C. and D. Jefferson (1994). Artificial Life as a Tool for Biological Inquiry. Artificial Life 1(1): 1-13.

- Taylor, T. (1998). Using Bottom-up Models to Investigate the Evolution of Life: Steps Towards an Improved Methodology. In C. L. Nehaniv and G. P. Wagner (Eds.), *The Right Stuff: Appropriate Mathematics for Evolutionary and Developmental Biology*, Workshop at ALIFE IV, UCLA.
- Taylor, T. (1999). *From Artificial Evolution to Artificial Life*, PhD Thesis. University of Edinburgh
- Thornton, S. (1997). Karl Popper. *Stanford Encyclopedia of Philosophy*
<http://plato.stanford.edu/entries/popper/>.
- Tinbergen, N. (1976). Derived Activities: Their Causation, Biological Significance, Origin and Emancipation During Evolution. *Quarterly Review of Biology* 27: 1-32.
- Toulmin, S. E. (1958). *The Uses of Argument*. Cambridge, Cambridge University Press.
- Trask, R. L. (1996). *Historical Linguistics*. London, Arnold.
- Trudgill, P. (1995). *Sociolinguistics: An Introduction to Language and Society*. London, Penguin.
- Trudgill, P. and J. Cheshire, Eds. (1998). *The Sociolinguistics Reader, Volume 1: Multilingualism and Variation*. London, Arnold.
- Werner, G. M. and M. G. Dyer (1991). Evolution of Communication in Artificial Organisms. In *Artificial Life II*. C. G. Langton, C. Taylor, J. D. Farmer and S. Rasmussen (Eds.), Addison-Wesley.
- Werner, G. M. and M. G. Dyer (1993). *Evolution of Herding Behavior in Artificial Animals*. In J. A. Meyer, H. L. Roitblat and S. W. Wilson (Eds.), *From Animals to Animats 2: proceedings of the Second International Conference on Simulation of Adaptive Behaviour*, MIT Press: 478-485.
- Whorf, B. L., Ed. (1956). *Language, Thought and Reality: Selected Writings of Benjamin Lee Whorf*. Cambridge, MA., MIT Press.
- Wooldridge, M. and N. R. Jennings (1995). *Intelligent Agents*. Berlin, Springer-Verlag.
- Worden, R. (1998). The Evolution of Language from Social Intelligence. In *Approaches to the Evolution of Language*. J. R. Hurford, M. Studdert-Kennedy and C. Knight (Eds.), Cambridge University Press: 148-166.
- Zahavi, A. (1979). Ritualisation and the Evolution of Movement Signals. *Behaviour* 72: 77-81.