

Language as an Exaptation:
Simulating the Origin of Syntax

Mark Bartlett

Doctor of Philosophy

University of York

Department of Computer Science

October 2006

Abstract

Much of the recent computational research into the evolution of language has concentrated on explaining the origins of compositionality and syntax in language. In such models, the ability of syntax to allow generalisation leads to it naturally emerging in the resulting language to capture structural properties of the semantic space under discussion. However, while such models can explain why protolanguages may have gained in structural complexity to become fully-fledged languages given the opportunity, they do not explain how the ability of individuals to handle composition of linguistic fragments evolved: while existing models may explain the emergence of syntax in the language, they presuppose a syntax-handling capability in the brain. It is the evolution of this capability that this research seeks to address.

This thesis lays out one possible explanation for the evolution of this linguistic ability and develops from it a computational model to assess its feasibility. Specifically, the biologically plausible idea is examined that the ability to handle compositionality in language is derived from a similar, and earlier, ability to handle compositionality in navigation and that the same underlying neural mechanisms are used. A second, supplemental, theory is also proposed, that one of the original purposes of language may have been for use in navigation. Communication in this case would be a form of inherently cooperative social behaviour which could lead to evolutionary benefits for groups of individuals possessing this trait. To assess the ability of these theories to explain the evolution of the capability of individuals to handle compositional language, a multi-agent simulation is created in which populations of agents with a variety of linguistic and foraging policies are tested for their abilities to survive and reproduce.

In addition to using the model to essay the relative successes of these behaviours, the role of the environment structure in determining the benefit of a behaviour is also examined. By varying parameters of the landscape, it is established that populations able to communicate can grow faster and be more resilient to volatility of resources than those unable to do so. In some situations, communication even provides the

ABSTRACT

ability for a population to survive when resources are too scarce or volatile for a non-communicating population to do so. Such results point towards a possible source of evolutionary pressure for the ability to use language.

Contents

1	Introduction	13
1.1	Nativism vs. Empiricism	15
1.2	The Origin of the Language Faculty	17
1.3	Outline of Thesis	22
2	Navigation	25
2.1	Localisation	26
2.2	Navigational Methods	27
2.2.1	Dead Reckoning	27
2.2.2	Beacons and Landmarks	29
2.2.3	Environmental Shape	31
2.2.4	Compasses	33
2.3	Cognitive Maps and Routes	34
2.3.1	Routes	35
2.3.2	Cognitive Maps	36
2.3.3	Evidence of Cognitive Maps	38
2.3.4	Routes of Landmarks	41
2.4	Summary	43
3	Neurological Evidence of the Basis of Syntax	44
3.1	The Declarative/Procedural Model	44
3.2	Broca's Aphasia	45
3.3	Language Use and Spatial Perception	48
3.4	Statistical and Rule-Based Learning	51
3.5	Summary	55

4	Computer Models of Language Evolution	57
4.1	Genetic Evolution	59
4.1.1	Lexical Communication	63
4.1.2	Compositional Communication	71
4.2	Cultural Evolution - Iterated Learning	75
4.2.1	Lexical Communication	77
4.2.2	Compositional Communication	81
4.3	Cultural Evolution - Negotiation	85
4.3.1	Lexical Communication	86
4.3.2	Compositional Communication	88
4.4	Summary	91
5	Using Artificial Life to Model the Evolution of Language	93
5.1	The Computer Model	94
5.2	Knowledge Representation	97
5.2.1	Aboriginal Songlines	97
5.2.2	Navigation	98
5.2.3	Language	100
5.3	Agents' Behaviour	102
5.4	Experimental Details	107
5.4.1	The Environment	108
5.4.2	Navigation and Communication	109
5.5	Model Viability	111
5.6	Theoretic Aspects of the Language	114
5.7	Summary	116
6	The Role of the Environment	118

6.1	Environmental Factors Assessed	119
6.1.1	Resource Spacing	120
6.1.2	Resource Volatility	122
6.2	The Effect of the Environment on Benefits	124
6.2.1	Resource Spacing	125
6.2.2	Resource Volatility	128
6.3	The Effect of the Environment on Relative Benefits	132
6.3.1	Resource Spacing	134
6.3.2	Resource Volatility	139
6.4	Summary	145
7	Comparison of Different Behaviour Types	147
7.1	Navigational Behaviour	148
7.2	Communicative Behaviour	152
7.3	Experimental Results	154
7.4	Summary	163
8	Examination of the Emerged Knowledge	168
8.1	Different Classes of Knowledge Utilisation	168
8.1.1	Knowledge as an Internal Phenomenon	169
8.1.2	Knowledge as an External Phenomenon	176
8.2	Language as an Altruistic Act	182
8.2.1	Knowledge as an Internal Phenomenon	183
8.2.2	Knowledge as an External Phenomenon	185
8.3	Summary	187
9	Conclusions	189

CONTENTS

9.1	Detailed Summary of the Thesis	191
9.1.1	Literature Review	191
9.1.2	Simulation Design	192
9.1.3	Experimental Results	194
9.2	Future Directions	196
9.3	Summary	199

List of Figures

1	The difference between cognitive maps and routes.	37
2	The mazes used to test cognitive mapping in rats (taken from Tolman et al. (1946)).	40
3	Results of the verbal shadowing experiment of Hermer-Vasquez et al. (1999).	50
4	The simple recurrent network (SRN), based on Elman (1990).	54
5	Modified SRN capable of functioning in two different domains, based on Dienes et al. (1999).	54
6	Classification of social behaviour by its effects on the fitness of individuals involved.	61
7	An example of a cell from the model of Ackley and Littman (1994). . .	65
8	Pay-off matrix for the prisoner's dilemma game.	74
9	The Expression/Induction cycle.	76
10	A meaning tree for a single channel, as used by Smith (2005)	80
11	Example of a grammar that may be found near the beginning of an experiment by Kirby (2002a).	83
12	Example of a grammar that may be found half-way through an experiment by Kirby (2002a).	83
13	Example of a grammar that may be found at the end of an experiment by Kirby (2002a).	84
14	An example of combining two exemplars from Batali (2002) to produce a new phrase.	90
15	A range of environments with different spacing of resources.	95
16	The behaviour followed by agents in the model.	104
17	Examples of the forward-biased movement of agents.	107
18	The configuration of landmarks used in experiments.	109
19	Environment for establishing the validity of the model.	112
20	Results of the experiment using the environment shown in figure 19.	113

LIST OF FIGURES

21	Typical environment layout for the resource spacing experiment.	121
22	Various environments used for resource spacing experiments.	122
23	The environment layout used for volatility experiments.	123
24	The role of resource spacing in determining population size.	126
25	Comparison of population dynamics for different values of d_{same} and d_{diff}	127
26	The role of volatility in determining population size.	129
27	Comparison of population dynamics for different parameters of volatility.	131
28	The relative benefit of language use over selfish behaviour as spacing varies.	135
29	The relative benefit of language use over non-linguistic altruism as spacing varies.	135
30	The relative benefit of non-linguistic altruism over selfish behaviour as spacing varies.	136
31	Comparison of populations implementing various behaviour types over time.	138
32	The relative benefit of language use over selfish behaviour in a volatile environment.	140
33	The relative benefit of non-linguistic altruism over selfish behaviour in a volatile environment.	140
34	The relative benefit of language use over non-linguistic altruism in a volatile environment.	141
35	Comparison of populations implementing various behaviour types over time.	143
36	The behaviour followed by type-1 agents.	149
37	The behaviour followed by type-2 agents.	151
38	Examples of Grammars	170
39	Distribution of number of rules per path used by agents.	179
40	Distribution of length of paths generated by agents.	181
41	The number of path generation attempts that failed.	182

42 Distribution of number of rules used and length of paths produced by
agents. 186

List of Tables

1	Results of Cheng’s experiment (Cheng, 1986).	32
2	Properties of routes and maps. Taken from O’Keefe and Nadel (1978)	38
3	Taxonomy of simulation models, showing examples of research published in each area.	58
4	Forms of navigation used in the model.	149
5	Forms of communication used in the model.	153
6	Average sizes for populations implementing each behaviour in the first environment.	160
7	Relative improvement between populations implementing each adjacent behavioural pair in the first environment.	160
8	P-values for comparisons of all behaviours in the first environment.	161
9	Average sizes for populations implementing each behaviour in the second environment.	161
10	Relative improvement between populations implementing each adjacent behavioural pair in the second environment.	162
11	P-values for comparisons of all behaviours in the second environment.	162
12	Average sizes for populations implementing each behaviour in the third environment.	164
13	Relative improvement between populations implementing each adjacent behavioural pair in the third environment.	164
14	P-values for comparisons of all behaviours in the third environment.	165
15	Average sizes for populations implementing each behaviour in the fourth environment.	165
16	Relative improvement between populations implementing each adjacent behavioural pair in the fourth environment.	166
17	p-values for comparisons of all behaviours in the fourth environment.	166
18	The number of paths known by different agent types	173
19	The length of paths known by different agent types	174

20	The number of paths known by different agent types	184
21	The length of paths known by different agent types	184

Declaration

I declare that all work appearing in this thesis is my own, unless otherwise acknowledged and credited in the text.

Some material upon which this thesis is based has been published elsewhere before;

- The computer model which is presented in chapter 5, and which is utilised in chapters 6 and 7, was originally presented in Kazakov and Bartlett (2002). It has subsequently been explained in varying degrees of detail in Kazakov and Bartlett (2004a,b, 2005) and Bartlett and Kazakov (2005a,b, 2006).
- The experimental results presented in section 6 are based on previously published experiments originally published in Bartlett and Kazakov (2005a) and Bartlett and Kazakov (2005b).

1 Introduction

Of the many unique attributes and skills of humans, perhaps the most important is not our capability for abstract thought or our ability to use tools, but our near-universal ability to use language. So important is this ability to both our everyday life and to our culture that to some extent it is part of what it means to be human. Without language, nearly all of humanity's largest ventures would have been impossible; how would it be possible to coordinate even a few dozen workers on a relatively simple construction project without language, never mind the building of Egypt's Pyramids or the great gothic cathedrals of Europe, or of a technologically complex device such as the spacecraft that took mankind to the moon. Without language, how would a social species such as ours have managed to create collective units bigger than tribes? How would the towns, cities and countries that have improved our security and prosperity been governed without access to language?

So central to our nature is language that in situations in which we find ourselves unable to communicate, we rapidly devise ways to do so. Enslaved people transported to the Caribbean often had no common tongue, but invented pidgins based on their own native languages and those of the sailors transporting them. Children born into these communities expanded and developed the pidgins into creoles. Elsewhere, groups of deaf children growing up without access to an existing sign language have invented their own system from scratch. Our desire to communicate is so driving that, where direct speech or gesture is impossible, we have invented methods to encode language in dozens of other media, including numerous written alphabets, hieroglyphs, pictograms, knots, electrical dots and dashes, and electronic binary codes.

Of all the features of human language, one that seems of particular significance is the ability to make use of syntax. Without this ability many of the other important features of human language, such as the ability to generate novel sentences or to include recursion, would be completely impossible. Indeed, in Jackendoff's analysis of the steps leading to language (Jackendoff, 2002), this ability is so fundamental that he considers

that even protolanguages must have possessed this feature.

While syntax is not entirely unique to human language, there are few other commonly cited examples of other animals possibly using it, and all in a way that is significantly different to how humans employ it. Amongst these species are birds, whose song is indeed syntactic: it is clearly made up of smaller units which are assembled according to given rules (Balaban, 1988). However, while syntax is employed in human language to allow compositionality, birdsong lacks a property that is integral to compositionality; the overall meaning of a birdsong is not linked to the meaning of its parts. Indeed there are no meanings associated with either the parts, or the song as a whole (except perhaps a message about the bird's fitness). Similarly, compositionality appears to be lacking in other communicating species who exploit syntax, including dolphins (Janik, 2000) and gibbons (Geissmann, 2002).

The only species other than humans who are known to have compositional communication are bees.¹ Bees employ a 'waggle dance' to communicate the distance and direction of sources of nectar to others (von Frisch, 1967). Part of this dance involves a movement, the duration of which indicates the distance and the angle of which indicates the direction. This movement, which combines two different meanings is often viewed as a form of compositionality. However, this composition of meaning is far from how compositionality functions in human language; the dance must convey exactly two meanings; no more, no less. Clearly bee dance expresses multiple meanings by a combination of movements, however the closed nature of the dance which prevents either meaning being expressed individually, or additional meanings being added, marks this as distinct from the compositionality of human language.

Given this almost complete lack of compositional syntax in communication systems other than human language, this thesis therefore seeks to answer two interconnected questions related to syntax. Firstly, how such an ability could have originally arisen

¹Recent evidence has suggested the presence of compositional communication in the warning cries of putty-nosed monkeys (Arnold and Zuberbuhler, 2006). However, the calls given do not exactly match in a one-to-one correspondence with the type of predator spotted or action taken, suggesting more data is needed to be certain true composition is being used.

in humans, and secondly, how it could spread through a population after it had been developed. As the first step in answering these questions, the remainder of this chapter will survey the major positions on the nature of linguistic ability in humans, and the possible origins of this, in order to clarify exactly how the views laid out in this thesis relate to current theory.

1.1 Nativism vs. Empiricism

Linguistic theories as to the ability of individuals to learn a language can be divided into two broad classes, *nativist* theories and *empiricist* theories. The former of these are the more established and are favoured by most of the more renowned linguists, though recent arguments from the empiricists have forced their position to be more seriously considered.

The originator and most prominent proponent of the nativist approach is Chomsky (Chomsky, 1965, 1980, 1995). Integral to his paradigm is the idea of dividing language into two parts: *I-Language* which is an individual's internal knowledge of a language, and *E-Language* which is the external behaviour they can exhibit as a result of their I-Language. It is to the study of I-Language, and in particular its acquisition, that nativists apply themselves.

The nativist theories are based on the observation that the acquisition of a fully grammatical language is a complex task, yet one that is accomplished with apparent ease by nearly all children exposed to one during a critical phase of development. From this, they hypothesise that there must be an innate language-specific capacity to learn grammar present in the brain, which they call *Universal Grammar* (UG). This UG specifies all possible grammars and any particular language can be selected from this grammatical space by a learner based on the linguistic input they receive. The UG can be seen as a series of parameters controlling all aspects of the grammar, allowing a particular grammar to be chosen by the setting of each of these parameters to a particular value; for example, in specifying English grammar, one would set the pre/postposition

parameter to the preposition setting, but would set it to the postposition setting as part of the specification of Japanese grammar. Chomsky suggests that there is a component of the brain that sets the UG parameters in response to linguistic input, which he calls a *Language Acquisition Device* (LAD). However, the inability to explain exactly what linguistic triggers are used to set the parameters, or indeed what parameters exist, is one of the major problems of the nativist approach.

One of the most powerful arguments for a nativist approach is the *poverty of the stimulus* argument (described by Pullum and Scholz (2002) amongst others). In brief, this argument states that the amount and type of linguistic input given to a child is insufficient to correctly choose the correct grammar from an unlimited space of possible grammars, therefore some innate bias towards the correct type of grammar must be present in the child (which is another way of describing the function of the LAD). Amongst the points raised by the poverty of the stimulus argument are the following:

- Only a finite number of sentences are presented to the child but the learned language will be capable of generating and understanding an infinite number of sentences, including ones not presented during the learning process.
- Only positive examples are ever given to the child: the child is not told that any sentences are ungrammatical.
- An infinite number of grammars exist which are consistent with any linguistic input.
- Only negative examples will allow a reduction in the size of the grammatical space, but are not given.
- Numerous uncorrected errors will be present in the linguistic input given, such as slips of the tongue.

In contrast to the nativist approach, empiricists believe that little or no language-specific abilities are needed to explain the acquisition of language: empiricists claim that the general-purpose capabilities of the brain are sufficient to explain most (if not

all) of the linguistic abilities, with the obvious exception of the neurological systems necessary to control the vocal cords, etc. Opinions of the empiricists range from the most extreme version of this position (i.e. that there is no LAD at all), such as that expressed by Sampson (1997), through to more moderate views such as those presented by Bates and Elman (1996), who propose a greater part of the grammar learning mechanism may be general-purpose than nativists traditionally accept.

Given that the poverty of the stimulus argument is the nativists' strongest backing, it is not surprising that empiricists are called upon to defend their position in respect of this evidence. While some of the points can be met with counter-arguments, such as the fact that errors in the input are unexpectedly rare (Newport et al., 1977), and that examples which allow certain constructions to be ruled out are present (Pullum and Scholz, 2002), some of the points are indisputable, such as the fact that an infinitely expressive grammar must be acquired from finite data. However, this does not present as great a problem for the empiricists as might be expected, as Elman (1993) has shown that general-purpose learning mechanisms are capable of performing this type of generalisation of linguistic data.

1.2 The Origin of the Language Faculty

If, deferring the empiricist position for the moment, the majority nativist view is subscribed to, one is forced to enquire as to the possible origins of the dedicated language faculty.² On this issue, opinions can be divided into three schools of thought. While all agree that innate grammatical abilities imply that the responsible faculty must be genetically specified, they disagree on how this may have arisen. One group of linguists hold that natural selection for linguistic ability must be the method that has led to the establishment of the language faculty; the second propose that the capabilities of the

²The term language faculty will henceforth be used to refer to the whole of the neurological systems dedicated to language, with the exception of those systems controlling input and output, for example the motor control needed to produce specific phonemes. UG and the LAD are just potential parts of this faculty.

language faculty originally evolved for other tasks and were then reused for linguistic processing; the final group suggest that no evolution of these faculties occurred and that they are a by-product of other evolutionary changes to the brain. These three groups could be labelled *adaptationists*, *exaptationists* and *architectualists* respectively.

In its purist form, the adaptationist position is that the language faculty is a specific neurological module that has evolved for the sole purpose of language use, with the selective pressure coming from an evolutionary benefit derived from the use of language: The ability to use language provides some specific reproductive or survival advantage, hence genes encoding the language faculty will increase the fitness of the organism they are found in. Pinker and Bloom (1990) espouse this position because they claim that;

Evolutionary theory offers clear criteria for when a trait should be attributed to natural selection: complex design for some function, and the absence of alternative processes capable of explaining such complexity. Human language meets this criterion: grammar is a complex mechanism tailored to the transmission of propositional structures through a serial interface.

One of the arguments frequently employed against this adaptationist position is to question how such a complex structure could have evolved. There are two potential answers to this question, either a *macro-mutation* (a single genetic mutation which creates the whole complex faculty in one step), or a series of smaller incremental mutations, each of which provide a benefit over the previous system. Both of these solutions strain the credibility of the anti-adaptationists; it seems unlikely that a single mutation should be capable of creating a complex system such as the language faculty, while an incremental evolution of the faculty forces one to ask what is the use of ‘half a grammar’.

While Pinker and Bloom (whose paper mounts the strongest defence of the adaptationist position) themselves rule out the suggestion of a macro-mutation, they are able to provide examples of scenarios in which features they consider to be part of the ‘design’ of the language faculty would provide a selective advantage. For example, they show that recursion allows reference to objects more precisely and that this can have an effect

on survival; ‘*For example, it makes a big difference whether a far-off region is reached by taking the trail that is in front of the large tree or the trail that the large tree is in front of.*’ Jackendoff (2002) also offers an adaptationist perspective on the stages of evolution which led to human language. However, rather than concentrate on how each of these would provide selective benefit, he outlines the order in which they must have occurred. For example, he argues that the ‘*development of a phonological combinatorial system*’ is predicated on the ‘*use of an open, unlimited class of symbols*’, which itself relies on the ‘*use of symbols in a non-situation-specific fashion*’.

Like the adaptationists, the exaptationists believe that the language faculty evolved through natural selection. However, they claim that the selective pressure driving this evolution was not originally language use itself. Gould and Vrba (1982) have coined the term *exaptation* to refer to this evolutionary mechanism, whereby a feature of an organism that evolved by natural selection to serve one purpose is later reused for a different purpose.³ For example, there is little doubt that the penguin’s wings originally evolved for flight, however a change of lifestyle has led to them being recycled to function as flippers. Similarly, it has been suggested that insects’ wings may have originally evolved to capture heat, only later being turned to flight (Kingsolver and Koehl, 1985). The exaptationist position is that the component parts of the language faculty can likewise best be explained if they originally served some other purpose in the brain. For example, the language faculty must possess the ability to handle grammatical categories; such an ability may simply be a recycling of an earlier, more primitive, ability to categorise objects in the world into classes such as ‘prey’ or ‘predators’. The evolution of the language faculty from pre-existing neural modules, each previously evolved to serve some non-linguistic purpose, can then be explained in one of two ways; as a macro-mutation linking together these modules, or through the cumulative agglomeration of these modules.

³The term *pre-adaptation* is sometimes used in a similar (though not identical) context to exaptation, however it carries an unfortunate connotation of foresight, as if evolution had chosen to select for a feature before the benefits associated with its possession arose. For this reason, the term exaptation will be used in preference throughout.

While some linguists have taken the former position, such as Bickerton (1990) (though he seems more recently to have altered his opinion (Bickerton, 1998, 2000)), most (including, perhaps surprisingly, Chomsky (1988)) have preferred the latter. In this interpretation, the only difference between the exaptationist and adaptationist explanations of the evolution of the language faculty is how each component that was added to the faculty was created; in theory, the adaptationists believe that the component evolved from scratch for the purpose of language, while the exaptationists believe that recycling of existing structures occurred. In truth, there is not such a clear-cut dichotomy between the two positions. Most exaptationists would probably concede that the pre-existing modules were not just linked together to form a language faculty, but were subject to some subsequent change and selection for language. Likewise, most adaptationists would admit the possibility that the components that form the language faculty may have originally had some other function, but have been selected for their use in language so much that they are greatly different from these original structures. The strict adaptationist and exaptationist positions can thus be seen as extremes on a continuous scale. In between these positions, there is a whole range of opinions in which the closer to the exaptationist position a theory is, the more it emphasizes the original evolution of the components for non-linguistic purposes, and the more towards the adaptationist position, the more important the subsequent evolution under selective pressure for language is deemed to be.

In contrast to both the adaptationist and exaptationist accounts that emphasize evolution of the language faculty, the architecturists believe that the language faculty may not have been directly selected for by evolution. Piattelli-Palmarini (1989) expresses this anti-selectionist viewpoint best by stating that *'[e]ven if a trait is useful and actually enhances the life expectancy of individuals who possess it, this fact does not grant the inference that the trait is there because it is useful.'* In contrast, architecturist theories claim that the language faculty is an inevitable side-effect of architectural constraints operating on the brain. Gould and Lewontin (1979) call such structures *spandrels* in a term adopted from architecture. In architecture, the spandrel is the triangular structure formed by the intersection of arches at right-angles. They can be

so elaborately adorned and decorated that it might be tempting to see them as having been deliberately created in order to bear these designs. However in reality, they are an unavoidable by-product of, for example, mounting a dome on arches. As Gould and Lewontin say of the spandrels of Venice's San Marco basilica;

The design is so elaborate, harmonious and purposeful that we are tempted to view it as the starting point of any analysis, as the cause in some sense of the surrounding architecture. But this would invert the proper path of analysis. The system begins with an architectural constraint: the necessary four spandrels and their tapering triangular form. They provide a space in which the mosaicists worked; they set the quadripartite symmetry of the dome above.

Analogously, the architecturist position states, the language faculty may appear to be complex and especially well-suited to its role, but it could be an inevitable consequence of other architectural constraints affected by other neurological systems.

In the same way that it was possible to bridge the gap between adaptationist and exaptationist theories, it is also possible to construct theories encompassing both architecturist and non-architecturist positions. A way of illustrating this is to look at architectural spandrels. While they are inevitable given the layout of the arches, they are not inevitably highly decorated; this is a decision that has been made by an architect or a craftsman. Likewise, the language faculty may have occurred as a result of the way the brain is formed, but again, it may have been possible for it to have been much simpler than it is. In the same way that a craftsman has taken the simple spandrel and turned it into an ornate display, natural selection may have taken a simple language faculty originating as a spandrel and, through evolution, turned it into a much more complex system.

Perhaps the best way to end this discussion, and to show how close the positions are in the debate over the origin of the language faculty, is with the following quote from Pinker and Bloom (1990) who, despite arguing strongly in favour of the adaptationist

position, are forced to admit the integral role that exaptation and spandrels may play;

Spandrels, exaptations, laws of growth, and so on can explain the basic plans, parts, and materials that natural selection works with – as Jacob(1977) put it, nature is a tinkerer, not an engineer with a clean drawing board. [...] But, as Darwin stressed, when such parts and patterns are modified and combined into complex biological machines fulfilling some delicate function, these subsequent modifications and arrangements must be explained by natural selection.

1.3 Outline of Thesis

The remainder of this thesis will focus on the evolution of a particular capability essential to human linguistic ability, specifically the ability to handle syntax. This can be divided into two major tasks, explaining the possible origin of this capability, and investigating how such a trait could be promoted and spread through a population. These objectives can be considered separately; if future evidence falsifies either part of the thesis, the other part may still be considered as a viable theory: parts of the model exploring the spread of the capability may still be valid even if the ability has a different genesis from the one outlined, while the evidence supporting this origin is not entirely dependent on the model developed.

The first, and most substantial, of the objectives will be accomplished by outlining how certain navigational abilities are functionally equivalent to parsing abilities, and building from this a theory that the linguistic ability could be a reuse of the earlier navigational capability: the central theme of this thesis will be to establish that syntactic language may have evolved from earlier navigational abilities. This can be viewed as either an empiricist theory, which suggests that a general-purpose ability is being used by language, or as a nativist theory that is most in accordance with the exaptationist opinion. As has been noted above, the theories of language faculty evolution are not completely disjoint however, and so given that the ability under discussion is so

fundamental to any language, this theory is not entirely antithetical to most nativist theories: an adaptationist could view the navigational abilities as a beginning point from which linguistic abilities could evolve, or an architecturist could consider the navigational abilities as forming part of the constraints that lead to the language faculty, in the same way that arches lead to spandrels.

The second objective, that of exploring the spread of the trait, is accomplished through computer modelling. A model is developed which allows the performance of populations of agents following different navigational and linguistic policies to be analysed.

The rest of this thesis will therefore proceed as follows:

Chapter 2 The major ideas from the literature on navigation are presented. Concepts such as allocentric and geocentric coordinates, dead reckoning, and beacons and landmarks are explained. In addition, the navigational abilities of various species including humans are explored.

Chapter 3 The emphasis then shifts to the grounding of language and navigation in the brain and, in particular, to the evidence of the link between linguistic and navigational abilities in the brain. Current work in neuroscience is presented to establish this link. One issue that is explored in particular is the separation of syntax and lexicon in the brain, an important factor as the simulation developed later will assume that syntactic abilities can be treated independently of lexicon formation, in contrast to most existing computer models.

Chapter 4 After the establishment of a link between navigation and communication in the previous chapters, this chapter moves onto the second objective of explaining the spread of linguistic abilities, by exploring recent simulations of language evolution. In particular, prominence is given to what these models establish about the origins of syntax.

Chapter 5 From the work of the previous chapters, a computer model is developed that can be used to explore the population dynamics of groups of agents implementing particular linguistic and navigational policies. The link between language and navigation is formalised here, in the representation used simultaneously by agents for both tasks. The general methodology for the performance of experiments is devised and explained.

Chapter 6 This chapter contains experiments exploring the role of the environment in the benefits of syntactical communication to a population. Specifically, simulations are carried out using the methodology of the previous chapter which show how the advantages of language use may vary depending on environmental factors. The factors studied are resource availability and volatility.

Chapter 7 Further experiments using the model developed in chapter 5 are presented in this chapter. The focus in these experiments is to determine a plausible route by which syntactic language may have evolved from an initial state of non-communication. To this end, a variety of navigational and communicative strategies are evaluated for the relative benefits they provide.

Chapter 8 After the model has been developed and been used to establish the plausibility of the theory advanced, the language which emerges during runs is analysed in its own right. Doing so provides evidence which further explains the results observed in the previous chapters.

Chapter 9 Finally, a discussion will draw the work together, reiterating and summarising the major findings of the thesis, and explaining how future work may potentially draw on the methodology and results.

2 Navigation

While the behaviour of agents in the experimental model presented in this thesis does not replicate precisely a form of navigation used in nature, it is nevertheless informed by concepts of navigation explored in cognitive science. Therefore, this chapter will present relevant research from this subject area, with a view to illustrating pertinent concepts. Particular prominence will be given to navigation by landmarks and the use of routes, which are both used in the computer model developed later.

In examining navigation, the process of locating one position from another, three reasonably distinct scales can be identified. The first of these is navigation over very long distances, such as that necessary for migration in birds. Such navigation is known as *large scale navigation*. At the opposite end of the scale, *small scale navigation* is concerned with locating exact positions in a local environment, such as for the relocation of previously buried food. Activity between these levels is *mid-scale navigation*, and is that used by most creatures in everyday situations, such as foraging for food.

Navigation at each of these different scales involves separate challenges. For example, birds migrating over oceans between continents may need to retain an established heading for periods of several days in the absence of guiding landmarks. In contrast, a squirrel recovering buried nuts will be likely to have access to several visual clues, but must be able to quickly integrate positioning information from each to retrieve its hoard.

High-quality scientific studies of navigation have been conducted for the better part of a century, during which time the prevailing theories and paradigms have changed substantially. For this reason, while citations in this section refer to the original published research, the interpretations given to the findings owe a great deal to secondary sources. These sources, while not presenting any original research, bring together and interpret experiments in the light of theories or findings unknown at the time the original research was published. Of particular use in understanding this large and evolving field were Shettleworth (1998) and Pearce (1997).

Section 2.1 now goes on to explore two classes of positioning, allocentric and egocentric localisation, which is followed by an outline of the major navigational methods that utilise these techniques in section 2.2. Finally, the important distinction between the use of routes and maps in recording and recalling navigational information is considered in section 2.3.

2.1 Localisation

In order to successfully navigate, individuals must possess the ability to *localise* items or places, remembering their position with respect to some frame of reference. In the most trivial case, this ability is necessary to return to a starting location, such as a nest, after foraging for food. Two contrasting methods of localisation exist, *egocentric* localisation and *allocentric* localisation, which differ in the frame of reference used.

Egocentric localisation is the term used to refer to remembering environmental positions relative to oneself. While such a technique is computationally simple, and therefore even possible by a creature as uncomplicated as an ant (Wehner, 1992), several disadvantages to the use of this method exist. Foremost among these is the accumulation of errors that occur. Through small discrepancies that affect an individual's movement, such as being gradually blown off course, inaccuracies in the expected position of objects accrue over time. Large numbers of turns, or turning through large angles, may further compound this error (Etienne et al., 1988; Maurer and Seguinot, 1995).

In contrast to egocentric localisation, allocentric (or geocentric) localisation operates by recording positions relative to external environmental clues, such as landmarks or magnetic fields. This technique overcomes the major limitation of egocentric localisation, that of accumulating error in an individual's position. By utilising a frame of reference external to an individual, continual readjustment of its assumed position can occur through comparison with the perception of the environment. However, successfully calculating one's own position relative to an externally defined frame of reference is in itself a problem. Many solutions to this issue have evolved in nature, depending

on the environmental clues used to establish position, and are explored with reference to the method of navigation they permit in section 2.2.

2.2 Navigational Methods

Given the large range of distances over which animals must travel, and their differing capabilities and purposes in making these journeys, it is not surprising that many different forms of navigation have evolved. However, a common goal is necessary in all tasks, that of building an internal representation of certain points in the environment for later recall.⁴

2.2.1 Dead Reckoning

For species that use egocentric localisation, the only method of navigation possible is that of *dead reckoning*. Dead reckoning allows an individual to track the relative position of an object, for example its nest, by continually monitoring its own movement and the direction of that movement. From this information, the individual is able to calculate the relative position of that object, and can return to it directly, without the need to return along its original path. It is important to note that the course taken in heading directly towards the object is not arrived at by perceptual properties that allow the detection of the target, but by an internally stored expectation of its location. As dead reckoning is limited to tracking a location relative to oneself, it is most commonly used by creatures to maintain a record of relatively few locations such as that of their home. It has been shown that dead reckoning is used in at least some circumstances by honey bees, to track the distance and direction to a food source for communication

⁴In some species, navigational information is stored externally through the use of chemical trails in the environment, the most commonly cited example of this being the pheromone trails used by ant colonies (Denebourg et al., 1983). However, as the overall goal of this thesis is to explore the link between the internal representation and processing of navigation and language, the forms of navigation examined shall be restricted to those used by species with internal representations. A review of the use of algorithms based on the behaviour of ants can be found in Dorigo et al. (1999).

in their waggle dance (von Frisch, 1967).

Experiments conducted on animals which use dead reckoning reveal problems with this method of navigation. When a desert ant is manually displaced before commencing its homeward journey, the path it subsequently takes is equal in direction and magnitude to the one it should have taken to return home from the position from which it was moved, rather than that from its new position to its nest (Wehner, 1992).⁵ After travelling this path, the ant subsequently searches for its nest using a spiral pattern centred on the location at which it expected to find the nest. Throughout this subsequent search, the ant will continually restart from the location at which it believed the nest was most likely to be, demonstrating that it continues to track this position through dead reckoning. Clearly, this distant physical displacement of individuals rarely naturally happens in the way described. However, flying creatures may find themselves slowly displaced from their expected location by the wind, and insects may likewise be moved by a strong gust of wind or a fast flowing stream of water.

Experiments with geese (Paul, 1982) and wasps (Ugolini, 1987) reveal that these species possess more sophisticated systems for determining their movement, and can track manual displacement through the movement of their visual field. Hamsters and gerbils also use additional abilities as an aid to dead reckoning in the form of the vestibular system (Etienne et al., 1988; Mittelstaedt and Mittelstaedt, 1982). This mammalian structure senses acceleration, and can be used to detect turning by external agencies. However, neither of these mechanisms provide a complete solution to the problem of disturbance by an external force. The visual system used by geese is only of use when the geese are able to see the passage of their surroundings, presumably one reason why the nocturnal gerbils do not use this system. The vestibular system that the gerbils do use is only sensitive to reasonably fast movement, and can be rendered useless by gradual motion.

⁵Experiments suggest that the ant's ability to correctly determine the 'correct' direction of travel, even when rotated during displacement, is due to the use of a sun compass, as described below (Wehner and Lanfranconi, 1981). More tentative evidence suggests that the distances may be recalled through the use of a 'step counter' (Wittlinger et al., 2006).

2.2.2 Beacons and Landmarks

One simple method of locating an object is to use perceptual properties of the object itself in order to locate and move towards it. The object in question is said to act as a *beacon* (or proximal cue). In addition to using visual characteristics of the object, smell and sound also provide useful cues to an object's location. While such navigation is undemanding enough for use by most creatures, those species limited to this form of locating objects would be restricted to movement in a small area around their home. In practice, the ability to recognise beacons may form only part of a species' navigational aptitude, with other mechanisms used to travel over longer distances; Collett (1996) identifies at least four forms of navigation present even in bees and wasps.

In contrast to the proximal cues formed by beacons, *landmarks* are known as distal cues, that is, cues away from the target location. When beacons cannot be used to guide individuals to their goal, the position of that goal relative to landmarks can sometimes be remembered and used to inform navigation instead.

In the laboratory, the seminal experiment demonstrating the use of beacons and landmarks is Morris' Water Maze (Morris, 1981). A rat is placed into a high-sided circular tank of water containing only a plastic cylinder, which provides a dry platform just above the surface of the water. As rats dislike swimming, the rat will attempt to reach the only dry surface it can, the cylinder. If the cylinder is painted black and the top is slightly elevated above water level, then it forms a visual beacon that the rat may use to reach a dry place. Rats learn to approach the cylinder, regardless of position within the tank, after few trials. In a second setting, the cylinder is made transparent, lowered to very slightly below the water level and the water is coloured white (through the addition of a substance such as milk or chalk dust). If the cylinder is placed consistently in the same location, rats placed within this tank also learn to find the dry surface, but based on visible cues outside the tank which provide a guide to the cylinder's location by acting as landmarks. When the cylinder is placed in a novel location, rats are still able to locate the black cylinder immediately, demonstrating that it acts as a beacon. In contrast, they repeatedly search for the clear cylinder in the location to which they

have become habituated, indicating that the external landmarks are being used. Experiments continually reveal that most creatures, and vertebrates in particular, are able to use different forms of navigation depending on the nature of environment.

A further experiment illustrating the use of landmarks is presented by Tinbergen (1951). In this study, he constructed a ring of pine cones around the entrance of a digger wasp's burrow, and allowed the wasp to become accustomed to this structure during a number of visits to the burrow. Following this, while the wasp was away from the scene, the pine cones were each moved a short distance in the same direction, so that the opening to the burrow was now slightly outside the ring of cones. Evidence that the wasp was using the ring as a landmark was obtained through its subsequent behaviour; wasps most often landed and searched within the ring for the entrance to the hole. The use of displacement has established that many creatures navigate using landmarks in this way, including rodents, birds and insects (Cheng, 1986; Cheng and Sherry, 1992; Cartwright and Collett, 1983). Further similar experiments have shown that closer and larger landmarks are used preferentially, both by bees (Cheng et al., 1987), and by jays (Bennett, 1993).

In Tinbergen's experiment, a ring of cones were needed to establish the location of the entrance: a single cone would have provided insufficient information to locate the burrow on its own. In general, it is the case that at least two landmarks (or a landmark with multiple features) must be present in order to unambiguously identify a position. Navigation using multiple landmarks can be achieved by remembering the displacement of each one from the location of interest.

Remembering explicitly the displacement of several landmarks for each location that an animal needs to recall would entail significant cognitive and memory requirements, especially for a creature such as the digger wasp studied by Tinbergen, which simultaneously tends offspring located in at least three burrows. Evidence instead indicates the use of 'retinal snapshots', at least in invertebrates, which record the pattern of stimulation on the retina (the appearance) when at the target location. By moving so as to reduce the discrepancy between the stored snapshot and the current view, an animal

closes on its target. Cartwright and Collett (1983) demonstrate that when the scale of the learned landmarks is altered, the distance at which trained bees search varies as though they were trying to match the scaled landmark with the memorised one. This finding has also been replicated in wood ants (Graham et al., 2004). In gerbils, a similar mechanism appears to function, though other factors also play a role (Collett et al., 1986; Goodale et al., 1990).

The retinal patterns needed for this are acquired by creatures when they leave a place of interest. When bees and wasps return to a location, they tend to approach from a consistent direction suggesting that the patterns they store refer to the landmarks when viewed from a particular orientation. However, both appear to generate the images needed by circling the location before leaving it. It has been proposed that moving around the location in this way helps to distinguish nearby landmarks from the background detail (Zeil et al., 1996); closer points appear to move more slowly, while distant ones blur across the retina.

Beugnon and Campan (1989) extend the concept of snapshots to hypothesise a ‘sequential file memory’. This memory would store a succession of snapshots arranged in temporal order. By moving so as to match its view with each snapshot in turn, a creature could use this method to retrace a previously travelled route between two given locations.

2.2.3 Environmental Shape

In addition to the use of landmarks in the environment to establish location, animals have also been observed to utilise the shape of the environment itself for this purpose. Cheng (1986) placed a rat into a rectangular box with two sides conspicuously longer than the others and with each corner marked with obvious visual and olfactory cues. After being allowed to eat food from a particular corner, the rat was removed from the box and placed into a second identical box, with food buried in the corresponding corner. Digging in a corner was taken as evidence that the rat believed food to be buried there. In one experiment, food was presented in a random corner in each trial,

Experimental Setting	Correct Corner	Rotational Error	Other Error
Random Corner	47%	31%	22%
Same Corner	76%	22%	2%
Sensory cues Removed	47%	53%	0%

Table 1: Results of Cheng’s experiment (Cheng, 1986). Figures show the corners searched by rats for each experiment. In the first experiment, food is placed in a different corner for each training episode. In the second, food is placed in the same corner for all training sessions. In the final experiment, the sensory cues nearest the food and in the geometrically identical corner were removed. See text for further experimental details.

in another it was consistently presented in the same corner, and in the final, the cues were removed from the corner at which food was presented and its geometric equivalent. The results of these experiments are shown in table 1.

Results from the first experiment show that almost half the trials resulted in the rat correctly identifying the corner, but with a significant number searching the corner which appeared geometrically identical. When consistently fed in the same corner, the number searching that corner increases, though a sizeable number of trials still resulted in the rat exploring the opposite corner. These results indicate a significant reliance on environment shape even when cues are available; though the ability to detect and use the cues cannot be doubted given the results of the second experiment. The final experiment demonstrates that in the absence of cues in the relevant corners, the rats are unable to distinguish between the correct corner and its geometrical identical partner, despite disambiguating cues in the other two corners. Cheng interprets these results as evidence of separate models existing for the processing of geometric and non-geometric information, with the third experiment illustrating the non-geometric model only being used to incorporate cues from geometrically important areas. Findings reported in which a similar experiment was conducted with young chicks (Vallortigara et al., 1990) reveal that a preference for the use of landmarks, though geometric information could still be utilised when landmarks were absent.

2.2.4 Compasses

Over very long distant travel, guidance by the use of landmarks is likely to prove impossible. Similarly, in situations where landmarks may be scarce, a system unreliant on them may be necessary. In these circumstances, *compasses* may be used. A compass is a mechanism for calculating a direction, based on the perception of an external stimulus with known direction. Three compasses have been confirmed in nature, the magnetic compass, the sun compass and the star compass, though doubtless others remain to be identified. In addition to these compasses, other sources of information may provide compass cues such as polarized light and infrasound.

While the continual movement of the sun makes it unsuitable for use as a landmark, an animal can navigate by the position of the sun, providing it possesses knowledge of how the sun moves and an internal clock. Given that the sun's position in the sky depends on both the time of day and the time of year, this form of navigation may appear infeasibly complex, however it is actually used extensively in the animal kingdom in species as simple as ants (Wehner and Lanfranconi, 1981) and bees (Dyer and Dickinson, 1996). Likewise, nocturnal creatures, particularly those that migrate, may make use of the corresponding star compass (Wiltschko and Wiltschko, 1991), though familiarisation with the constellations in early life appears essential in its use.

The use of magnetic compasses is most associated with birds, despite the fact that it remains a disputed topic. While evidence exists that pigeons are sensitive to magnetic fields (Gould, 1982; Walcott, 1978), attempts to cause them to react to this stimulus in the laboratory have proved largely unsuccessful. In contrast, honey bees and tuna have both been unequivocally shown to possess this ability. Nevertheless, it appears likely that some birds do indeed possess magnetic compasses but that sun compasses are their preferred form of navigation. Only when the sun is unobservable do they switch to their magnetic sense.

The use of a sun compass in animals can be illustrated by training them in a spatial task at a given time of day and monitoring their performance in a subsequent test situation.

In this test, the animal's internal clock is first altered by holding it indoors for several days, with the room lit and darkened artificially so as to shift the apparent sunrise and sunset by a set number of hours. Following this, the animal is reassessed at the original task carried out at the original time. If an animal is relying on landmark-based navigation, its performance at the original task should be equivalent to that obtained during the training phase. However, should the sun compass be the primary form of navigation, the path taken by the animal should be altered due its shifted internal clock.

In an experiment of this type, Wilkschko and Balda (1989) trained scrub jays to locate food at a given location in an outdoor setting. While navigation by both landmarks and sun compass were possible, the birds ignored the landmarks in the test situation, and searched in the (incorrect) location suggested by the use of the sun compass. A similar experiment conducted on chickadees (Sherry and Duff, 1996) demonstrated comparable results, though intriguingly these birds only relied on their sun compass when tested in the presence of recognised landmarks, despite ignoring these landmarks in their navigation. In bass, the use of a sun compass was demonstrated more simply through their inability to navigate in cloudy conditions (Loyacano and Chappell, 1977). On clear days, bass which were moved from one end of a pond into the centre of a circular pool were able to return to the side of the pool corresponding to the end of the pond from which they were taken; on cloudy days, they moved to an apparently random edge.

2.3 Cognitive Maps and Routes

Most of the above discussion focussed on the methods by which animals may retain a memory of a particular site, but the issue of how travel between these locations may be achieved is also important. In many situations, travel can occur over long distances or in terrain that prevents the observation of beacons or landmarks from distance. In these circumstances, not only must a method of recognising the target location exist, but also the ability to travel between locations. For example, the digger wasp was

shown to locate its burrows by the use of landmarks near to the burrow, but as these landmarks would not have been visible from the distant destinations visited by the wasp; some other mechanism must have been available to guide the wasp to the general vicinity of the burrow.

One method of retaining the information necessary to make this journey has been mentioned in the form of dead reckoning. Using this method an individual can return to a previously memorised location by recalling the distance to the location and the direction in which to travel. More general, however, is the issue of how individuals with knowledge of many locations may retain the information to travel between them. Two possible answers to this problem exist, *routes* and *cognitive maps*.

2.3.1 Routes

Routes are the more primitive and less contentious form of representation. At its most basic level, a route consists of a method for travelling from one point to another based on a list of instructions to be followed in order. By comparison with a guidebook of walks for ramblers, O’Keefe and Nadel (1978) show that these instructions each consist of stimulus–response–stimulus commands: each instruction starts in a location identifiable by a specific stimulus, and specifies the action to taken in order to end in another recognisable location.

This formulation of routes is such that they specify some sequence of actions such that they lead to a particular outcome, immediately implying intention associated with each route. The goal-driven nature of routes, in conjunction with their low information content and the encoding of responses as part of the rule, inevitably allows high-speed processing of the path to take. This fast recall and utilisation of routes is their greatest advantage.

However, the simple nature of routes, as well as leading to their rapid handling, is also their major shortcoming. As routes specify a series of actions to take in specific situations, they prove very fragile to disruption. If a location along a route is altered, a

creature following that route will become confused and unable to successfully complete the path (Reese, 1989). Likewise, the route becomes useless for the purpose of navigation if the creature is moved from its known route, or blockages forcing a detour are encountered. Once away from a route, the creature is lost, which may occur for many reasons, ‘*The route instruction can be physically or mentally degraded [...] landmarks or other guides may be destroyed or changed [...] the translation of the instruction into behaviour may go awry [...] a momentary lapse of attention leads to a subgoal being missed.*’ (O’Keefe and Nadel, 1978). In each case, it is the stimulus–response–stimulus nature of the route which renders it useless: once a required stimulus is absent or the incorrect response is made, the remainder of the route becomes unfollowable. Navigation along the route can only be re-established through randomly searching until a recognised location is discovered.

2.3.2 Cognitive Maps

In contrast to routes, which describe paths between two specific locations, cognitive maps store the relative position of all known locations within a single structure. A consistent definition of a cognitive map does not exist, with researchers choosing to define it differently, but the major aspect which is widely agreed on is that it consists of a representation of the world, in allocentric space, which stores the location of known objects.⁶ A cognitive map may reasonably accurately be said to be a neural equivalent of an aerial view of an environment or of the type of map to be found in an atlas.

In contrast, if the knowledge of an individual navigating by routes were to be plotted as a map of the environment, the resulting diagram would resemble the kind of map used by train companies; each location would correspond to an intersection, and routes between them to lines (see figure 1). However, information on distances or directions between locations would not necessarily be accurate (if given at all), and crucially, novel means of travelling from one location to another without journeying along the routes

⁶Gallistel (1990) supports a broader definition of a cognitive map which would admit any system allowing displacements of objects to be remembered, including by dead reckoning.

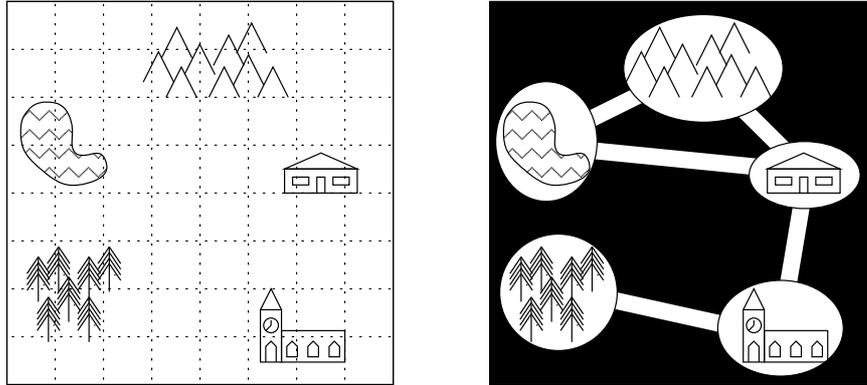


Figure 1: The difference between cognitive maps and routes. Using a map, one can plan a path directly between any pair of known locations. Using routes, one is restricted to following paths that are already known. For example, one can travel directly between the forest and the mountains, if using the map shown above on the left, whereas if using the routes shown on the right, this journey requires at least three stages, going via the church and house.

could not be generated.

It is the ability to generate novel paths which is the major advantage of maps over routes. The use of dead reckoning has been shown to allow a novel path back to the origin of a journey, however, in an animal with a cognitive map, journeys between any known locations should be possible. The ability to generate novel paths derives from two properties of a map. Firstly, maps represent locations as if on a system of allocentric coordinates. Given the spatial position of two locations in such a system, a route between them can be calculated. Secondly, whereas routes record a particular path to be taken to a target, maps contain only the relative positions of the locations in question. Should a route become blocked, or a creature deviate from a known route, the use of a map would allow an alternative route to the location to be generated based on the revised environmental knowledge.

While it is in the *ability* to generate new routes that the advantage of maps is found, it is the *process* of generating routes that is their major weakness. Routes store sequences of actions to be carried out in given associated circumstances allowing rapid deployment. When using a map, in contrast, a route must be generated from the map, probably necessitating decisions on which of several potential routes is better and requiring the

	Route	Map
Motivation	The final stimulus is the goal; the route is built with this in mind	No object or place on the map is the goal; the map is usually built out of curiosity
Flexibility	Routes are rather rigid; they are rendered useless by any damage, or by the loss of a guidance, or direction	Maps are extremely flexible, and relatively invulnerable to noise and damage
Speed	Very fast	Relatively slow
Information Content	Relatively little; each route contains only a small amount of data	Maps are one of the most efficient information storage devices known, with very large capacity
Access	No specialised knowledge is required for access; no coding strategies	Special knowledge of coding strategies required
Manipulation	None	Maps can be compared; places on maps can be compared

Table 2: Properties of routes and maps. Taken from O’Keefe and Nadel (1978)

relevant information in the map to be converted to a form from which a route can then be followed. This takes time and adds additional neural complexity to that needed to follow a route (O’Keefe and Nadel, 1978).

The relative merits of the use of routes and maps for navigation are shown in table 2.

2.3.3 Evidence of Cognitive Maps

The presence of cognitive maps in animals is a much debated issue. The primary method of experimentally testing for their existence is through the ability to generate novel routes between known locations. As has previously been shown, animals navigating through the use of routes are unable to generate novel paths between locations, whereas those navigating by map should possess this ability. Attention has mostly been placed on the capabilities of rats and other mammals, especially other rodents.

The most famous study involving the ability of bees to generate novel routes is that of Gould (1986). In this experiment, bees were first trained to fly along a route between their hive and a feeding site. The bees were then captured at the hive and transported to another site, chosen such that it was out of sight of the feeding site, and formed an equilateral triangle with the other vertices located at the hive and the feeding site. When released, the bees spent some time apparently orientating themselves before flying in the direction of the feeding site. The ability to travel from the release location directly to the feeding site is taken as evidence of cognitive mapping: the bees were able to observe the hive, and based on this infer their position and thus calculate a route directly to the feeding site.

Unfortunately, other explanations for the observed behaviour are also available. Dyer (1991) argued that the bees could have calculated the path to take based on the angle of travel relative to certain landmarks visible in the environment, however Gallistel (1994) argues that this would have required complex calculations to be performed. Alternatively, it is possible that bees had flown from between the point of release and the feeding site before, during the training phase; such behaviour is hard to observe or prevent. In this situation, the behaviour exhibited would correspond to following this known route. Finally, the view from the novel release position could have been sufficiently similar to the feeding site for snapshot memory to be used to allow the bees to navigate by reference to the location of visible landmarks.

Dyer (1994) performed a modified version of Gould's experiment in which the feeding site was placed in a quarry and the hive and novel release site on higher ground. The results obtained from this matched those of Gould. However, when the location of the feeding and release sites were switched, so that the release site lay in the quarry and the feeding site was on the higher level, the bees did not fly directly to the feeding site after release, but carried on travelling in the direction they were following when they were caught. Dyer uses these results to cast doubt on the conclusions made by Gould. Dyer suggested that the bees in his first experiment flew up, were able to obtain a view similar to that from the hive, and were thus able to navigate towards the food. When the bees began in the quarry, their initial orienting flight would not have allowed them

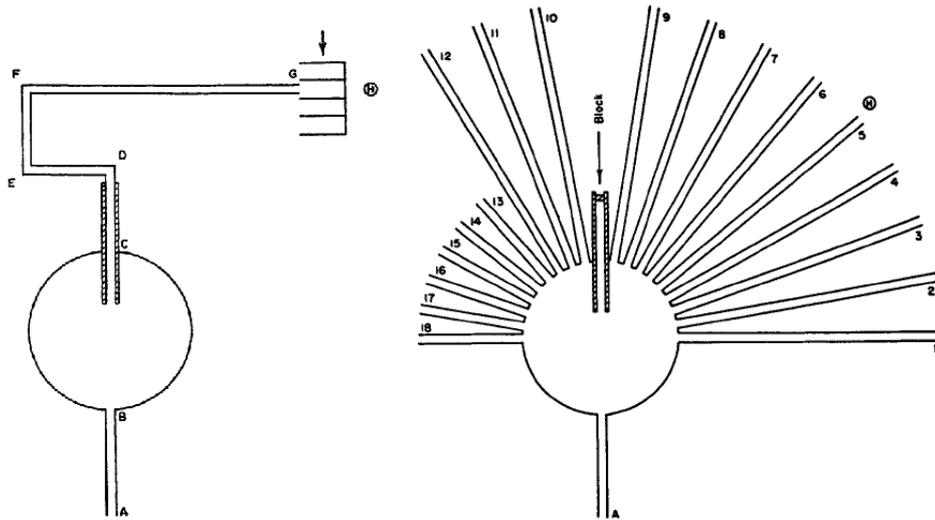


Figure 2: The mazes used to test cognitive mapping in rats (taken from Tolman et al. (1946)). The maze on the left shows the training configuration and that on the right the test situation. In each case the rat is released into the circular arena from the bottom. In the training situation, food was placed at the end of the route, as indicated by the arrow. In the test situation, most rats choose to run down arm 6, towards the position of the food in training. The circled 'H' at the top-right shows the location of a lamp.

to obtain a view similar to that obtained from the hive, and hence they were unable to travel to the feeding site. Dyer's evidence points to the use of navigation over large distances based on snapshots rather than a full cognitive map as Gould suggests.

An early experiment to test whether cognitive maps were possessed by rats was carried out by Tolman et al. (1946). In this experiment, a rat was trained to run to a particular position in the environment in order to obtain food (see first diagram in figure 2). The environment was subsequently altered except for the original position of the rat, and the geographical position of the food (second diagram in figure 2). Rats were still able to locate the food, which was taken by the experimenters as evidence that they possessed a map allowing them to understand the geographical position of the food relative to the starting position: rats possess a cognitive map.

As with the experiments in bees however, the results are open to interpretation. In the environment, a light was placed above the location of the food. By associating this bulb with the food, a rat in the test phase could have used this bulb as a beacon in

order to travel to the food. Should this have been the case, no need for a cognitive map would have been shown to exist.

In an experiment involving hamsters, Chapula and Scardigli (1993) used a maze shaped like a six-spoked wheel to demonstrate the use of a cognitive map through its ability to allow detours to be calculated. As explained earlier, a route-based system of navigation falls apart should the route become blocked. A map-based system however, should allow the calculation of a new route incorporating the fact that a blockage now exists. The hamsters were initially trained to collect food from a location in the environment by blocking all passages except those on a route which led between the start location and the food. Following the training phase, all the blockages were removed and another created along the route that the hamster had been taught. The testing revealed hamsters to alter their behaviour to most often take the new shortest route to the food. This result is consistent with the idea that the hamsters possess cognitive maps which recorded the position of the start and target locations.

2.3.4 Routes of Landmarks

Up to this point, it has been assumed that creatures are able to create routes without mention of what form these may take. While most work on routes has been with the aim of establishing whether behaviour can be explained in this way, or if it requires access to a cognitive map, some research on the nature of routes suggests that, in some species at least, routes can be represented through the use of a sequence of landmarks.

It is known that bees do not always travel by the most direct path between locations, but head towards landmarks situated alongside their route. Chittka and Geiger (1995) investigated the role that these landmarks play, by positioning large artificial landmarks between a bee hive and a feeding location. Once bees had become accustomed to the original configuration, the order of these landmarks was altered. Subsequently some bees were observed to land prematurely after passing the landmark originally placed closest to the hive. When the original configuration of the final two landmarks was presented further from the hive, an even greater number landed in the wrong location.

These results demonstrate that not only did the landmarks play a role in the definition of the route, but that bees possess some knowledge of the concept of sequence.

Further evidence of the use of landmarks in routes is obtained from experiments on ants (Collett et al., 1992). A similar setting to that used above with bees was created by placing two landmarks between a feeding site and the nest. One landmark was placed to the left of the ant's path, and the other to the right. When these landmarks were moved, after the ant had become familiar with their position, the path followed by the ant on later foraging trips was modified to include a detour to ensure that each landmark was passed on the side to which the ant had become accustomed.

A final experiment illustrating this point, was conducted using butterfly fish (Reese, 1989). Butterfly fish are a species living amongst coral reefs, and follow a particular path between feeding sites and home. When part of the coral along this route was moved, the fish halted their normal journey at the point at which they expected to find the missing coral. Further experimentation revealed that fish moved from their known route moved randomly until they were able to locate a familiar section of the route, at which point they resumed their normal behaviour. The later part of this experiment has also been replicated in wasps (Baerends, 1941).

While each of these observations is consistent with the use of routes, it is worth noting that they do not rule out the use of maps. In each case, the behaviour witnessed is most easily explained through the creatures involved possessing knowledge of routes formed of sequences of landmarks. However, it is possible to reconcile the results of the studies with the use of a cognitive map. If one assumes that the species studied possess cognitive maps, the results can be explained through individuals revising their supposed position in the environment based on the landmarks they observe. In the final experiment, the random movement undertaken before reaching a previously utilised route would make sense in a creature possessing a cognitive map if that creature had been repositioned to part of the environment it had not previously explored and therefore mapped.

2.4 Summary

This chapter has explored major ideas from cognitive science surrounding the concepts of navigation. In particular, various methods of navigation have been outlined, and the differences between navigation by map or by routes have been made clear. As far as possible, this survey focussed on accepted results in this field and has avoided cutting edge or contentious issues, save an examination of the possibility of cognitive maps in animals.

In examining the issue of the existence of cognitive maps in animals, an attempt has been made to present the relevant research in as neutral manner as possible, without preferring map-based or route-based explanations for experimentally observed behaviour. In part this derives from the fact that the experimental work in this thesis has nothing to add to either side of the debate and in part is due to the computer simulation that is developed in later chapters utilising the uncontroversial method of routes for navigation.

Evidence of route-finding in nature using a sequence of landmarks was presented, due to its similarity to the mechanism utilised in the model developed in chapter 5. This method of navigating underlies the fundamental observation of this thesis, that particular aspects of navigation are functionally equivalent to processes in the production and comprehension of language. In the following chapter, this link is further explored by the investigation of the neural structures behind navigation and language.

3 Neurological Evidence of the Basis of Syntax

In the introductory chapter, the twin issues of the extent to which language required specific learning capabilities and structures, and the evolutionary origins of these systems were examined. This chapter revisits these issues with the aim of investigating the neurological systems underlying the language faculty. Specifically, the intention is not to reiterate the earlier debate, but to establish which cognitive abilities share neural structures with linguistic processing components.

Various approaches to this are taken. In the following section, a model of the brain, the declarative/procedural model, which seeks to explain the structures underlying linguistic ability is presented. This model highlights known systems which may play a role in the processing of language. Following this, evidence of the position within the brain of language structures, based on studies of patients with brain lesions, is presented. Section 3.3 presents further experimental evidence suggesting a common neurological basis for the processing of language and navigational abilities. Finally, selected studies showing the ability of animals, children and artificial neural networks to acquire different classes of language are presented in order to explore the extent to which grammar-processing abilities may be shared with other species.

3.1 The Declarative/Procedural Model

Linguistic knowledge consists of two broadly separate types, lexical and grammatical knowledge. Noting this, Ullman (2004) presents a model of neuro-linguistic embodiment called the *declarative/procedural model*. In this model, the functions of the mental lexicon are associated with the declarative memory system (Eichenbaum, 2000), while grammatical abilities are taken to have commonality with the procedural memory system (Mishkin et al., 1984).

The declarative memory is concerned with ‘*the storage and use of knowledge and facts*’ (Ullman, 2004). The system, which is situated in the temporal lobe, has been highly studied and is known to be important in allowing fast learning and plays a crucial role in the formation and retrieval of short-term memories. Importantly for a possible role in the mental lexicon, it is involved in the learning of associative information, such as that which is needed to link semantics, morphemes and phonemes.

In contrast, the role of the procedural memory is to learn and process stimulus-response rules and to administer sequences of such rules, ‘*whether the sequences are serial or abstract, or sensori-motor or cognitive*’ (Ullman, 2004). It is this sequence processing ability that implicates the procedural memory system in underlying grammatical competence. While grammatical structure is more complex than simple sequential rules, hierarchical processing abilities are associated with this system, at least in humans. Procedural memory is located primarily in basal-ganglia and frontal structures, including pre-motor regions and Broca’s area.

Situating lexical and grammatical knowledge in particular neural systems in this way allows specific predictions to be made about linguistic ability based on understandings gained of those other abilities associated with these structures. This is especially true of the declarative memory system, which is more studied and better understood than the procedural memory. For example, the effect of oestrogen in improving declarative memory has been well studied (Maki and Resnick, 2000; Miles et al., 1998; Sherwin, 1998). Should the declarative memory also be involved in the mental lexicon, similar findings correlating oestrogen and linguistic ability should also be made. Some initial evidence of this may be found in women’s superior declarative memory use in verbal memory tasks (Kimura, 1999).

3.2 Broca’s Aphasia

Aphasia is the name given to various classes of linguistic disorder resulting from damage to the brain. Given knowledge of the area of the brain that has been impaired, study

of aphasic patients offers a potential tool for matching diminished cognitive abilities to the locations where these capabilities are grounded in the brain. Of particular interest to this thesis is a type of aphasia, *Broca's aphasia*, which results from damage to the region known as Broca's area. Patients with this type of aphasia exhibit difficulties in producing and understanding syntactic or morphological aspects of language.

Should the abilities underlying language production and comprehension utilise neural structures involved in other capabilities, one would expect to see a decline in both capabilities if the structure involved were to be damaged. Likewise, should an individual be trained to improve their performance on a task involving the damaged area, the improvement would be expected to transfer over onto other abilities that also utilised the same area.

Based upon the latter of these two observations, Hoen et al. (2003) construct a study in which they measure the effect of training Broca's aphasics to re-order non-linguistic sequences on their ability to understand sentences in which the natural word order of active sentences has been altered to form more complex structures such as passive or relativised sentences. Active sentences are the most readily understood sentence structure, in which agent, action and object occur in that order,⁷ for example 'The cat chased the ball'. Comprehension of more complex sentence types (for example the passive 'The ball was chased by the cat' or relativised 'It was the ball that the cat chased') requires transformations of the elements into the simpler active sentence order. If the elements in these sentences are numbered, then transformation from one sentence type to another can be expressed as a reordering rule, for example $123 \rightarrow 321$ captures the restructuring needed to form a passive sentence from an active one.

Hoen *et al.* utilise the fact that reordering rules of the above type can apply equally to linguistic and non-linguistic components to train aphasics the required transformation to form an active sentence from a relativised one using symbolic playing cards. The experimental setting used involved an initial test phase, to establish the inherent abilities

⁷This is true of English, which is a subject-verb-object language. Other languages may follow different orderings for their most common sentence types, such as subject-object-verb in Japanese.

of their subjects, training on non-linguistic re-ordering across a period of ten days and then a test phase to measure any change in their subjects' abilities. Comparison of the tests conducted at the beginning and end of their training period reveals that patients performed equally well on understanding active and passive sentences before and after training, but crucially improved significantly their ability to comprehend relativised sentences after the training phase.

The simplest coherent explanation of the observed result is that the sequence learning conducted with non-linguistic re-ordering was made available in the task requiring linguistic transformation: some neural structures involved in learning the re-ordering task formed part of a more general module capable of handling not only motor skills, but linguistic ones as well. However, evidence contradicting the conclusions of this study is to be found in an experiment by Goschke *et al.* (2001) in which skills learnt by aphasics in a motor task do not transfer to a similar task involving linguistic skills.

In contrast to the learning procedure used in Hoen *et al.*'s task, in which patients were aware of being taught a particular rule, Goschke *et al.* trained subjects using *implicit learning*, where subjects are oblivious to the fact that they are acquiring information subconsciously. The standard method of measuring subconscious acquisition of sequentially related data is Nissen and Bullemer's *serial reaction task* (SRT) (Nissen and Bullemer, 1987). In this framework, an asterisk appears at one of four locations on a computer screen and participants must press a key corresponding to the location as quickly as possible. Unknown to the subject, for sections of the trial a relatively short repeating sequence of positions is used. During the presentation of this sequence response times should decrease if the participant has learnt the pattern either explicitly or implicitly. Subjects' rate of performance often increases in these sections, even if they are unable to later express the pattern after the experiment has concluded.

In an experiment involving the standard SRT, Goschke *et al.* were able to demonstrate achievement by Broca's aphasics which was slower than for a control group (who were presented with the same stimulus but did not suffer from Broca's aphasia), but otherwise in line with the execution that would be expected from subjects capable of

absorbing the repeating sequential pattern. However, in a modified version of the SRT, in which rather than locate the asterisk, subjects had to indicate the position on screen of a letter which they had heard spoken, differences were observed between aphasics who responded by pressing a key and those who responded verbally. In the case of subjects who pressed keys, as with the original SRT, response times decreased when the repeated sequence was presented. However, aphasics who were required to respond verbally showed no improvement in the task during the period of the repeated sequence. An experiment with a control group showed standard patterns of response time for both manual and verbal response tasks.

The conclusion that should be drawn from this experiment in isolation is that sequence processing in motor control and that involved in language is based on different underlying neural circuitry. This is however, in opposition to the findings of Hoen *et al.*'s experiment. One possible way to reconcile the findings may be through the different forms of learning that were used in the experiments. In the second experiment, the aphasics were unable to implicitly learn a sequence that was to be used verbally. In the first experiment however, this portion of the experiment had no analogue. The sequence re-ordering was carefully taught to the subjects over a period of days allowing them to become proficient at it. If a similar approach had been taken in the second experiment (that is to say if the subjects had been allowed to familiarise themselves with the sequence non-verbally and then attempted to transfer this skill) perhaps the observed results would have differed.

3.3 Language Use and Spatial Perception

Further evidence linking the mechanisms underlying language use and other cognitive tasks is provided by Hermer-Vasquez *et al.* (1999). Whereas other studies highlight the similarity between sequence-based tasks and linguistic syntactic capability, this finds a connection between those navigational abilities needed to integrate different forms of spatial information and the use of language.

The study is carried out using a variation of the experiment of Cheng (1986) (presented previously in section 2.2.3) in which rats were placed in a rectangular box, and their ability to utilise both cues and geometry was assessed. It is unnecessary to repeat Cheng's findings in full again here, but to reiterate his major conclusion, rats were more sensitive to environment shape than they were to cues, particularly when those cues were away from the location of interest. In a version of Cheng's experiment conducted on human adults and children (Hermer and Spelke, 1996), it was found that children performed similarly to rats in failing to integrate geometric and other positional information. As with rats, preference was given to geometric information leading to rotational errors occurring. Adults correctly utilised both sources of information to locate the right position.

Hermer-Vazquez *et al.* utilised a similar setting to that used to assess the abilities of children to study the performance of adult humans attempting to locate a position while simultaneously engaged in another activity. In the experiment, a subject is placed in a white rectangular room with identical boxes in each corner. Two of the room's walls are shorter than the others and one of these is coloured blue. The subject observes the hiding of an object in a one of the boxes, is disorientated by rotating in the centre of the room with their eyes closed, and is then asked to find the hidden object. If the subject uses both the geometry of the room and the relative position of the coloured wall, then perfect retrieval should be observed. In contrast, utilising only the geometric information leads to both the correct location and the opposite corner being searched equally often. The use of only the colour cue leads to both corners next to or away from the coloured wall being searched (depending on the position of the object) and all corners will be searched with equal frequency if neither source of information is utilised.

In addition to a control group, the ability of individuals undertaking *verbal shadowing* was assessed. Verbal shadowing is a task consisting of simultaneously listening to a speaker and repeating what is being said. Compelling subjects in the experiment to verbally shadow recorded speech while observing and locating the position of the object forces the parts of the brain that play a role in language to be utilised for verbal shadowing, preventing them simultaneously processing data for any other purpose,

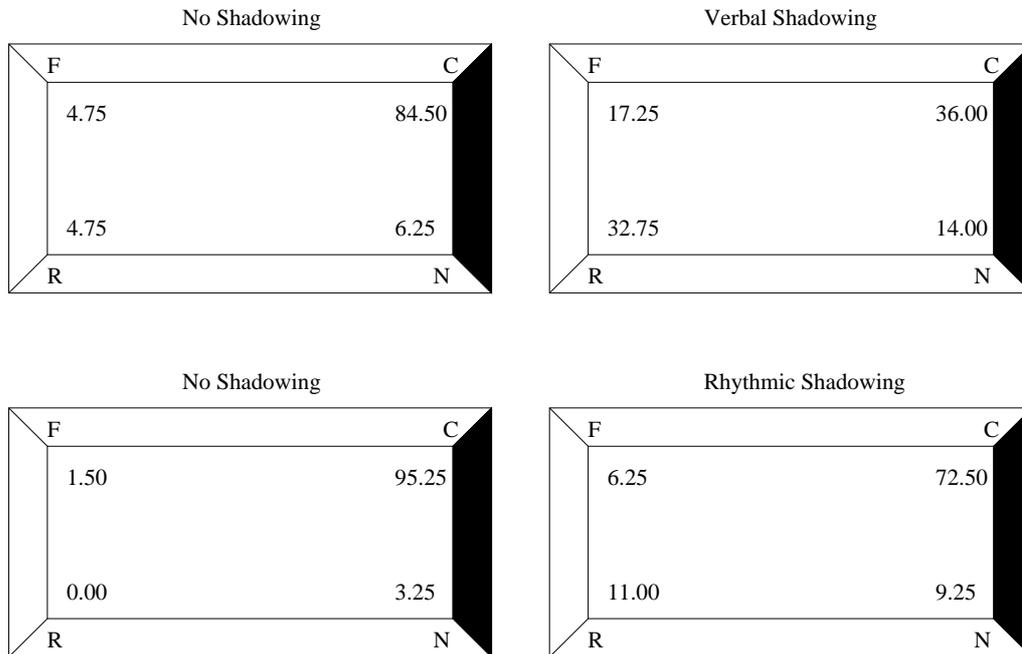


Figure 3: Results of the verbal shadowing experiment of Hermer-Vasquez et al. (1999). The top diagrams show the behaviour of the verbal shadowing group when either engaged in this task or not, and the lower diagrams show equivalent figures for the rhythmic shadowing group. In each case, corner C is the correct location and R is the geometrically identical corner. Numbers show the percentage of searches performed in each corner.

including the positioning task.

The results obtained from this experiment proved consistent with the hypothesis that verbal shadowing prevented the integration of geometric and cue information in the task (figure 3). Additional studies conducted using individuals engaging in a non-verbal ‘rhythmic shadowing’ appear to demonstrate that the inability to integrate the information is a result of the linguistic nature of the task and not the complexity of the verbal shadowing.

That the verbal shadowing interfered with the ability of subjects to combine various sources of information seems undeniable. However, two explanations for the cause of this seem possible. Firstly, language may have a role to play in remembering the location of the object; some evidence for this exists in the development of children’s cognitive abilities (Hermer-Vazquez, 1997). However, a more likely explanation is that some neural structures used in the integration of navigational information are also involved

in language processing tasks. In the experiments, the use of verbal shadowing would have forced these structures to take in their linguistic role leaving the navigational task with only the primary source of information, the geometric shape of the environment.

3.4 Statistical and Rule-Based Learning

In seeking to understand the grounding of language in the brain, it is necessary to establish the extent to which language, or parts of language, may be the result of domain-general abilities and the extent to which it requires its own unique processing systems. Of particular importance is the question of whether grammatical abilities, and those necessary to handle higher-order grammars especially, are exclusively available for language.

In order to study this point, comparison between three classes of learner may be carried out, in addition to the assessment of the abilities of adult humans. Studies on human children at a pre-grammatical phase of their linguistic development help to elucidate those aspects of grammatical acquisition that are innate: those abilities predefined in the brain which do not depend on experience of language to develop. Similar experiments on non-human animals can also be of use in determining which aspects of grammatical learning are uniquely human. Such studies are usually carried out with primates, as these creatures are most closely related to humans and therefore should exhibit behaviour most similar to our own; any difference in the abilities of humans and other primates is most likely to be explicable in terms of that behaviour having evolved in humans since the existence of our closest common ancestor. Finally, the behaviour of general-purpose learning algorithms can be investigated. In particular, *connectionist* models can be used to identify which aspects of language learning are possible using only statistical learning methods.

Studies of grammatical competence in primates and children are conducted using a common experimental setting. The structure followed is to first familiarise a subject with examples constructed according to a particular grammar and then to test their

ability to abstract the grammar by examining their response to novel sequences that are either consistent with the previously used grammar or that violate its principles. The initial familiarisation stage of these studies is conducted by simply exposing the individual to examples for a relatively short period of time.⁸ Testing in adults can be performed by asking the subject whether a test phrase was grammatical, however in children and animals, this is clearly not possible. Recognition in this case can be assessed by the period of time that the subject shows an interest in a test sample; individuals remain looking at a stimulus longer when exposed to a test sample they find unusual.

Of particular interest in this field is the question of whether algebraic rules, such as those needed to recognise context-free languages, are learnable by infants or by animals. A study of this type using cotton-top tamarins appears to show that they are able to discriminate successfully between grammatical and ungrammatical examples of the regular language $(AB)^n$, which consists of any number of repetitions of a pair of A and B syllables.⁹ However, they appeared unable to learn the simple context-free language A^nB^n , in which a number of A syllables are followed by the same number of B syllables, despite having n limited to 2 or 3 (Fitch and Hauser, 2004). Adult humans were able to learn either grammar. The authors interpret this finding as suggesting that the ability to learn the algebraic rules necessary for the more complex language type¹⁰ is unique amongst primates to humans.¹¹ In an earlier, similar study using seven-month-old infants (Marcus et al., 1999), a demonstration of the ability to learn algebraic rules was also claimed. In this study, infants were familiarised with either the grammar ABB

⁸Statistical learning as used here appears very similar to the implicit learning used by Goschke et al. (2001) earlier in this chapter, but significant differences do exist, as explained by Perruchet and Pacton (2006).

⁹Where A and B refer to categories containing mutually-exclusive sets of syllables throughout.

¹⁰While the more complex language is claimed to be context-free, it is strictly the case that as with any finitely-sized language, a regular grammar exists that is able to recognise it. However the size of the grammar is larger for this regular grammar than the equivalent context-free grammar.

¹¹Evidence exists that European starlings (and very likely many other species of songbird) are capable of learning to discriminate between regular and context-free grammars (Gentner et al., 2006), though almost certainly as a result of the ability evolving independently of humans, long after our last common ancestor.

or *ABA* and then examined on a test set. Crucially, the syllables used in the test set were different to those used in the familiarisation stage, so it is claimed that the infants must have learned the rules, not the relationship between particular syllables.

The explanations given for this finding has been called into question (Altmann and Dienes, 1999; McClelland and Plaut, 1999). Specifically, it has been queried whether the ability to learn algebraic rules is necessary to explain the findings of Marcus et al. (1999). It has been established that infants possess statistical learning capabilities, the ability to automatically recognise the relative frequency of transitions between syllables. It is claimed that these abilities are sufficient to learn the position of word boundaries using the familiarisation-preference procedure (Saffran et al., 1996), though their capacity to explain grammatical structures in which dependencies between non-adjacent syllables exist has been doubted (Pinker, 1997). By creating neural networks which implement statistical learning techniques, but are capable of discriminating between the languages in Marcus *et al.*'s study, Altmann and Dienes (1999) claim to have shown that grammatical learning may be possible without the ability to use algebraic rules.

The model used by Altmann and Dienes (1999) (which is described in more detail in Dienes et al. (1999)) is based on the simple recurrent network (SRN) developed by Elman (1990) (figure 4). A SRN consists of a three-layered neural network, with input, hidden and output units, which is used to model temporal sequences. At each time-step, the next symbol in the input is encoded at the input layer, and the predicted next element is read from the output layer. In order to provide the network with memory, an additional set of context units exist into which the contents of the hidden layer are copied at each cycle. These units can then be used as input to the hidden layer at the next time-step. Dienes *et al.*'s modification to this network involves the insertion of an additional layer of units, the encoding layer (figure 5). The insertion of an encoding layer allows the network to acquire a grammar based on one set of inputs and successfully recognise it even when the inputs used in the test phase differ from the training phase, as in Marcus *et al.*'s experiment.

McClelland and Plaut (1999) provide an explanation of why this works.

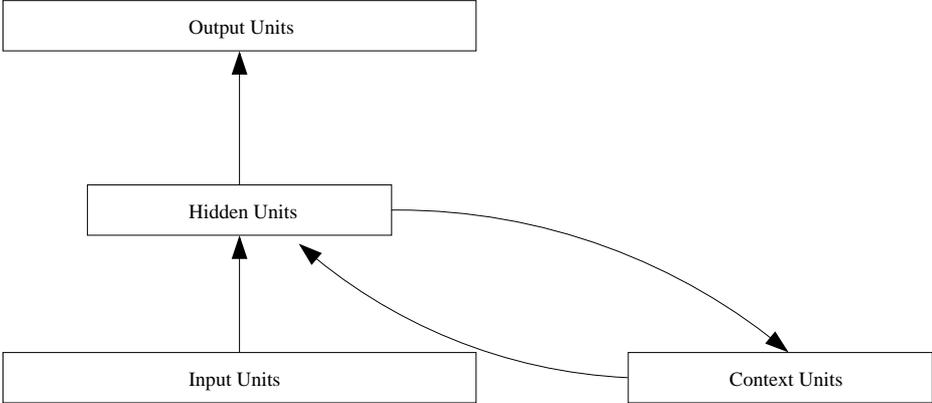


Figure 4: The simple recurrent network (SRN), based on Elman (1990).

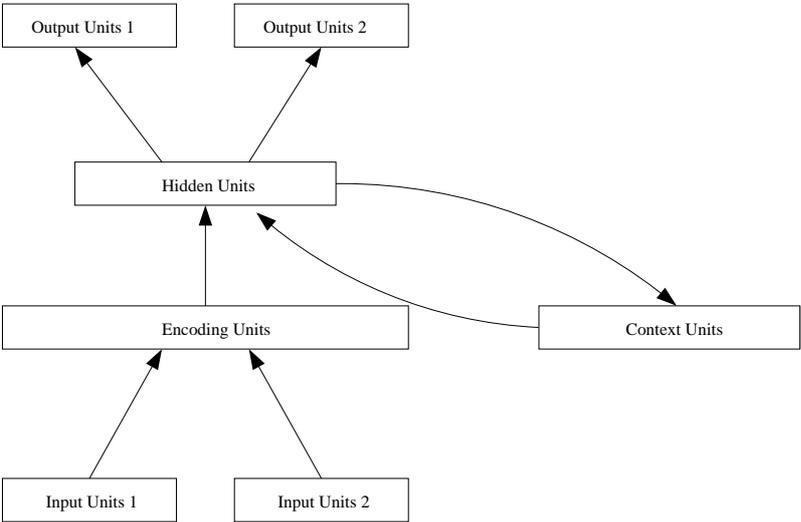


Figure 5: Modified SRN capable of functioning in two different domains, based on Dienes et al. (1999).

[F]or prior learning to generalise to a new stimulus, the representation of the new stimulus must overlap with [...] the representation of the stimuli on which learning is based. [...] [T]he characteristic of neural networks just described has been misconstrued as implying that the relevant overlap must be present in the input itself.

The relevant overlap of representations required for generalisation in a neural network [...] need not be present directly in the ‘raw input’ but can arise over internal representations that are subject to learning.

In alleging that the domain-general method of statistical learning may play a larger role in the learning of language than traditionally thought, Dienes *et al.* potentially extend the overlap between linguistic processing neuro-systems and those of more general-purpose computation. However, while the mechanisms may be the same, an innate predisposition may still exist towards using these capabilities for linguistic purposes (Clark *et al.*, 1997).

3.5 Summary

This chapter has presented results from various studies into the neurophysiological grounding of syntactic language capabilities in the brain. Additionally, the issue of what other capabilities may share the use of these structures has been addressed.

A model of the language capabilities in the brain was presented which drew a distinction between syntactic and lexical processing of language (Ullman, 2004). The syntactic capability, on which this thesis focusses, was claimed to be situated in a region of the brain that deals with sequence processing in many different domains. This link between sequence processing in linguistic and non-linguistic domains was further enhanced by experiments conducted by Hoen *et al.* (2003), though other experimental findings questioned this straight-forward link (Goschke *et al.*, 2001).

Additional research suggested a link between the processing of language and skills

needed for navigation (Hermer-Vasquez et al., 1999). Specifically, structures involved in language processing were demonstrated to play a role in integrating positional information from different sources. Later chapters of this thesis return to the link between navigation and syntax, though with the intention of demonstrating functional equivalence between route following and language parsing, a connection related to the processing of sequential information.

Finally, the uniquely human aspects of syntax were explored through comparison between the cognitive abilities of human adults, infants, animals and simple computer learners. While this field of work is still in a state of great debate, it seems clear that the statistical learning technique used by computer modellers in this area may be able to explain a greater proportion of the observed results than had traditionally been thought. The domain-general nature of this learning, along with its operation on sequential data, adds further weight to the argument for the components underpinning linguistic syntax capability to be positioned in the procedural memory system with other sequential processing abilities.

4 Computer Models of Language Evolution

The use of computer simulation in investigating the evolution of language is becoming relatively well established, having been utilised for nearly two decades, and over the last 5 years papers attempting to overview the literature have begun to appear.

Many of these have grouped work into those attempting to address the same issues in language evolution. For example Cangelosi and Parisi (2001) arranges studies into categories such as “Evolution of Syntax”, “Grounding of Language” and “Behavioural and Neural Factors”. Another categorisation similar to this is used by Nehaniv (2005), which focusses on the aspects of language evolution that yet remain to be adequately explained through simulation. Such surveys are useful in addressing the progress that has been made in pursuing solutions to specific research questions.

In contrast, a second, smaller set of survey papers organises research by the methodology used rather than its aims. One such paper is Wagner et al. (2003). In this review, studies are divided along two dimensions in to structured or unstructured communication and in to situated and unsituated agents.¹²

As there is currently no generally used categorisation of the simulation models, a taxonomy is proposed here which divides the research along two axes; the compositionality of the communication system under consideration and the method by which knowledge of that system is transferred from individual to individual. This taxonomy is shown, along with research typifying each area, in table 3.

A fundamental dichotomy exists along the first axis, between systems in which it is possible to combine signal/meaning pairs to produce utterances amalgamating the meanings of the original signals, and those in which it is not. The first group of commu-

¹²The frequently cited review paper by Kirby (2002b) sits somewhere between the two survey approaches in that its early chapters group papers with similar methodologies (e.g. “Innate Signalling Systems”) while its latter focus on the aims of studies (e.g. “Grounding and the Origins of Meanings”).

	Innate	Learned	
		Iterated Learning	Negotiation
Lexical	Cangelosi and Parisi (1998)	Oliphant and Batali (1997)	Steels (1996b)
Compositional	Hashimoto and Ikegami (1996)	Kirby (2002a)	Batali (2002)

Table 3: Taxonomy of simulation models, showing examples of research published in each area. The examples given are drawn from those mentioned in the text.

nication systems can be called *compositional* and the latter *lexical*, though they this second group also incorporates systems such as animal warning calls. Systems exhibiting morphology can be considered as a special type of compositional communication. This dimension corresponds closely to the structured/unstructured axis in Wagner et al. (2003).

Along the second axis, communication systems can again be divided in to two types; systems which are specified genetically, and evolve through natural selection, and systems which are learned by individuals during their lifetimes. These communication systems are called *innate* and *learned* respectively. Models involving learned transmission can be further subdivided into those involving lifelong learning by agents (*negotiation*) and those in which learning takes place just once for each agent, when it is created (*iterated learning*). The distinction between innate and acquired communication systems is not as marked as that between compositional and lexical communication, as it is possible to construct models in which individuals are created through genetic evolution and subsequently modified through phenotypical learning (MacLennan and Burghardt, 1993), indeed the interplay between learning and evolution has itself been the object of study (Smith, 2003b).

As the presented categorisation aims to categorise by the type of simulation used, it has a drawback in so much as it fails to capture the aspect of language evolution that simulations seek to address. For example Steels and Kaplan (1999) and Smith (2005) fall into different areas of the taxonomy, though both are concerned with the creation and diffusion of perceptual meanings. However, as simulations in each of the areas generally seek to address different problems in language evolution, and as the primary

aim of many experiments to date appears to have been either the establishment of coordinated lexical systems or the emergence of grammar, this shortcoming is perhaps not too important.

The remainder of this chapter now turns to a survey of the major existing models of language evolution following the structure of the taxonomy proposed. There are several reasons for embarking on this review. Firstly, to provide context to the novel experimentation presented in later chapters by illustrating the nature of existing simulations and the manner of their findings. In addition, the exploration of this previous work should prove that the performance of the experiments in this thesis seeks to answer a novel yet important question within this field. Finally, issues studied in other simulations, such as the origins of altruism or the emergence of a coordinated lexicon, prove relevant to the model developed in chapter 5.

4.1 Genetic Evolution

The genetic evolution of communication is concerned with the creation of coordinated systems through the action of natural selection on the speakers and listeners of the system: the transmission and reception behaviours are specified as part of the genetic make-up of the individuals, and any changes to them must come through the evolution of the population. Evolved systems such as these are found in nearly all species of animal which communicate and even in plants, in the use of colouration and scent of flowers to attract pollinators. These communication systems are distinct from those known as languages, which are learned communication systems in which some aspect of the system is acquired during an individual's lifetime.¹³

¹³It is often the case that the term *language* is used to refer to the fully-formed human language that exists today, with *proto-language* being applied to those learned systems that lack the full power of human language. However, as no computer simulation yet has demonstrated a communication system with complexity even remotely comparable to human language, the term *language* will be used in this thesis to refer to both full language and proto-language, unless it is necessary to distinguish between the two.

As human language is clearly learned, it may not immediately be apparent why the study of genetically encoded communication systems should be of any use in its understanding. However, the study of this type of communication, in addition to providing insight into animal communication systems, allows certain aspects of language to be studied more clearly through the removal of learned behaviour from models.

One such aspect is the evolution of the willingness to speak, which would be complicated by evolving what is generally taken to be a genetic propensity while assigning fitness to individuals based largely on their phenotypes (in the form of the language they had learned). While an advantage exists in being able to comprehend the speech of others, it is less immediately apparent what benefit is accrued through an individual producing speech itself. Clearly, some situations exist in which communicative accuracy is mutually beneficial, but many other situations exist in which signalling appears to occur while offering no advantage to the speaker. As Noble (1999) states;

Why do speakers freely exchange valuable information when the theory of natural selection predicts selfishness? In a hypothetical proto-linguistic community, what would prevent the rise of a selfish mutant strain that listened but did not speak? Speaking or signalling always costs something in terms of time and energy, and may involve more indirect costs such as exposing the signaller to greater predation risk. Why not reap the benefits of the informative signals of others, without paying the costs of signalling oneself? Or worse, why not use the communication system to lie, misinforming others for one's own benefit?

Communication acts are a particular case of social behaviour, and when considering any social behaviour, acts can be grouped into one of four categories depending on how they affect the expected fitness of the instigator of the act and the recipient. If an act tends to increase or not affect the fitness of both parties, it is called *mutually beneficial*. A common example of this type of behaviour includes pack hunting, which allows all participants to achieve a bigger meal than they would expect to receive as individuals. If the act has a positive or neutral effect on the performer, while having a negative effect

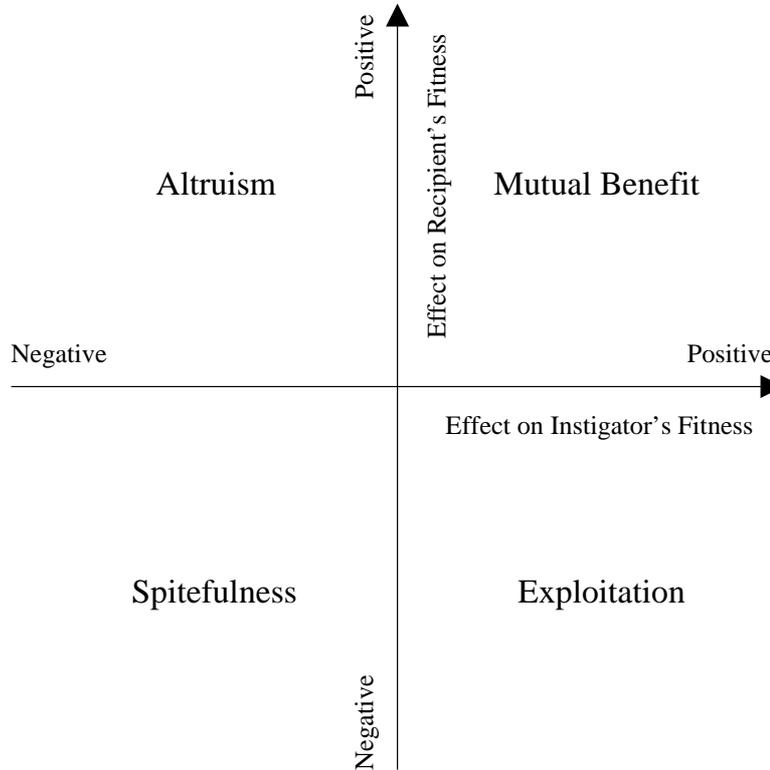


Figure 6: Classification of social behaviour by its effects on the fitness of individuals involved.

on the other individual then it is said to be *exploitation*, and occurs, for example, in the behaviour of the cuckoo in leaving its eggs in other birds' nests, resulting in the nest owners having to work hard to feed the child, rather than the cuckoo doing so itself. *Altruistic* behaviour consists of those actions in which the individual responsible for the act incurs a cost while the recipient of the action benefits. A common example of this in many species is child rearing; the offspring gain an enormous benefit from being feed and protected, while the parents place themselves in greater danger than usual and must forage more. A final class of action, *spitefulness*, occurs when an act is carried out with the deliberate intention of harming both individuals involved. However, such behaviour is rarely seen in nature, as evolution acts such to select against it. These classes of behaviour are shown in figure 6.

In situations in which communication is mutually beneficial or exploitative, it can be assumed that the willingness of an individual to speak will be promoted under

classical interpretations of evolution: in these situations, communication is expected to provide an increase in the fitness of the speaker by definition and so would be promoted through natural selection. However, these situations form only a subset of those in which communication occurs; in many situations in which individuals communicate, the exchange appears to be altruistic. As these acts decrease the fitness of the speaker, classical evolutionary theory predicts that individuals carrying out these acts will not persist in the population, yet observation of this behaviour in the real world shows that this is not so.

Several mechanisms have been proposed to account for the existence of coordinated communication systems in situations in which communication is an apparently altruistic act and will be outlined below in the context of simulations which demonstrate their occurrence. In each case, what on the surface appears to be altruistic behaviour is revealed to be self-interested behaviour when viewed from a different perspective.

Through the use of systems involving innate communication, a second factor can also be illustrated, namely that the benefits associated with communication are not solely related to the communication system an individual uses, but rather to the mutual comprehension between individuals: the similarity of the systems used by different individuals in the population. For example, while objectively, fully-formed modern language may be superior to the proto-language spoken during the early stages of language evolution, an individual from that time subject to a macro-mutation that gave him access to modern day grammar would have no advantage over his compatriots, as the extra expressiveness or disambiguation would not be understood by anyone he addressed. On the contrary, any evolved constraints that might help to structure the language may actually make it harder to comprehend the speech of a majority lacking this ability. More generally, access to a more powerful language does not necessarily equate with an increase in fitness; fitness will only be increased by improvements to the language that offer advantages in communicating with others. As selection must occur in models involving innate systems of communication, this method provides a better framework for study of this concept than models involving learned communication, in which such selection must be appended to the basic system.

In addition to the factors mentioned above, Cangelosi and Parisi (1998) also identify (without elaboration) ‘selective reproduction’, ‘changes in population variability’ and ‘drift’ as issues that may be studied with equal validity in both culturally evolving and genetically evolving languages.

4.1.1 Lexical Communication

Werner & Dyer The earliest paper to use computational modelling to study innate communication systems was Werner and Dyer (1992). In their model, communicative signals are produced by immobile females to guide blind males towards them in order to mate, which results in a pair of children being created and replacing two randomly selected existing agents. Despite reproduction occurring without an explicit fitness function, some agents are clearly fitter than others: those agents finding mates more regularly will tend to have more offspring than others. Both sexes are implemented as recurrent neural networks, with males mapping input signals to output movements, and females taking the position and heading of males as input and producing signals as output. Though initially females all produce different signals in order to guide males, and males all respond to signals differently, a communication protocol eventually emerges which is shared by all agents. Before this situation is reached, agents pass through the following stages;

- Random female signalling and male movement
- Extinction of males who stand still
- Evolution of males that always move in a straight line
- Evolution of a signal to tell males to turn left or right to find female
- Evolution of females to use left/right signals optimally

Through this model, the issue of an individual’s fitness depending on both its own communication system and those of others is illustrated more clearly than in perhaps any subsequent research. This can be observed through consideration of the fate, at

different points of the simulation, of a pair of males: one taken from the final population, and the other taken from near the beginning of the simulation.

At the end of the simulation, the male taken from the final population performs optimally: the females know how best to attract a male, and he can understand their signals. However, near the beginning of the simulation, he will do considerably worse; females will not yet have learned the sequence of behavioural steps to try to induce in a male to guide it, and the intention behind the signals sent will not correspond to the male's interpretation of them anyway. At an early enough stage of the experiment, this male will be outperformed by those who do not react to the confused and contradictory signals. By comparison with a non-communicating population, Werner and Dyer were able to establish that the ability to communicate actually reduces the fitness of individuals before the evolution of a coordinated turn signal, with females directing some males to spin in place due to the miscoordination of their communication systems. In contrast to the male from the final population, an average male from an earlier stage of the game will move in a straight line, the optimal policy for covering the maximum amount of squares in the shortest possible time, and the best policy possible without a communication system available. However, if placed in the later stages of the simulation, this male will perform considerably worse than the other males. By retaining the policy of ignoring the signals, the male will find fewer mates than the other males, who will be able to find mates not just in the row they travel, but also those in adjacent rows who signal them.

The much diminished fitness of these males when placed in populations at the wrong stage of development is a direct result of them having communication systems that are ill-suited to communicating with those around them. In particular, the outcome of placing an agent from the final stages of the simulation in an earlier population raises an important point: evolution will not inevitably act so as to promote what appears to be a more advanced communication system. While inspection by an outside observer may show the final system to be more advanced (allowing for unambiguous turning signals to be given), such an agent fares less well early in simulation than agents ignoring directions. In several simulations of the evolution of communication

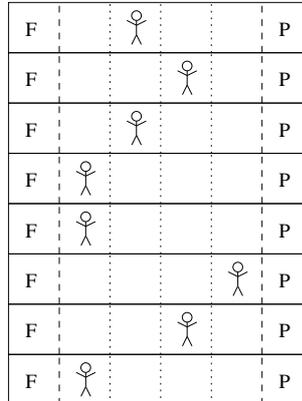


Figure 7: An example of a cell from the model of Ackley and Littman (1994). Eight individuals are shown in parallel tracks, with food (F) at one end and a predator (P) at the other.

(for example the model of Zuidema and Hogeweg (2000) discussed below), modellers create fitness functions that reward novelty of expression. The results produced by this model suggest that novelty will not necessarily be selected by evolution, and that such fitness functions should be carefully justified if the results of such experiments are to be accepted.

Ackley & Littman Ackley and Littman (1994) also develop an artificial life model to study the evolution of innate communication systems with the aim of investigating this evolution in situations in which successful communication evolves in the absence of benefit to the sender: they concern themselves with the problem of the establishment of a coordinated communication system in an altruistic population. In this model, each individual sits on a track, consisting of 4 squares, along which they can move. At either end of this track are locations which can contain either food or a predator. Individuals and their tracks are placed along side each other in groups of 8 to form cells, such that individuals can move along their own track, but are unable to move into or observe the state of other tracks. The positioning of predators and food is common to all tracks in a cell. The world is formed from an array of 128×128 cells. An example of one such cell is shown in figure 7.

A ‘day’ in the model consists of 36 trials, allowing each individual to experience every

possible combination of starting position and food and predator positioning. In each trial, all individuals are given a period of time to move along their track and then each receives a reward based on their individual final position (relative to any food or predators) and the number of movements they made. In addition to making a move at each time step, individuals can also choose to emit a signal on any number of six sound channels common to agents sharing a cell. At the subsequent time step, each individual can take into account these signals emitted by the others in the cell, as well as whether they are adjacent to food or a predator, in deciding which move to make. Producing or observing these signals has no direct cost for the agents involved. After each day, two of the most successful individuals in each cell are chosen to mate, with the child replacing a random member of that cell.

Ackley and Littman study three strategies to transfer behaviour from one cell to another, and find different outcomes result from each. In the first case, *wind*, after every 5 days, a single randomly chosen individual from each cell is ‘blown’ into a neighbouring cell. The second strategy is known as *festivals*. In this strategy, festival days occur every two days, and on them, instead of mating taking place within a cell at the end of the day, two of the most successful individuals from each group of 2×2 cells are chosen, and mate in the usual way. Finally, in the *windy festival* strategy, both windy days and festival days occur at set frequencies, every tenth day and every second day respectively.

When solely using wind to transfer individuals, the system converges towards a situation in which the majority of the population adopt the best strategy possible without communication. Within this population, cells in which communication is reliably used repeatedly develop, but are constantly destroyed by individuals who produce misleading signals. In the festivals only case, the world comes to be entirely dominated by altruistically communicating agents, but not before several selfish ‘species’ arise, including at least one that not only disrupts the communication of the cell it is in, but also acts as a ‘free rider’, exploiting the communication system to gain knowledge without contributing itself. When both wind and festivals are used, clusters of communicating agents form as in the festive only situation, but are eventually invaded and destroyed

by disruptive signallers, as in the wind only scenario.

In seeking to explain the observed behaviour, Ackley and Littman appeal to the notion of *kin selection*. This theory, which is also known as kinship-driven altruism or kin altruism, is closely related to the neo-Darwinian paradigm (Dawkins, 1976), based on the concept of inclusive fitness (Hamilton, 1964). Kin selection seeks to address the problem of why altruistic behaviour exists between relatives, in apparent contradiction to the theory of natural selection, by viewing the basic unit of selection not as the individual, but as the gene. When considered from this perspective, altruistic behaviour by individuals is consistent with natural selection when it is equivalent to selfish behaviour by the individuals' genes; altruistic acts towards relatives may be selfish from the perspective of a gene. As relatives share a proportion of genes that is inversely related to the degree of kinship between two individuals, a gene promoting altruistic action towards relatives will be selected for by evolution, provided that the decrease in fitness of the actor is exceeded by the increase in fitness of the beneficiary multiplied by the probability of the altruistic gene being present in this beneficiary. Using this insight, it is possible to explain behaviour such as child-rearing. When examined from the classical viewpoint, this behaviour consists of altruistic behaviour of parents towards children, however neo-Darwinism allows it to be seen as the genes in the parent improving the survival chances of those identical genes in the offspring: the parent is acting altruistically, because its genes are acting selfishly. Social colonies of ants can also be explained through kin altruism, though the situation here is complicated by the unusual genetic nature of these creatures, which results in siblings sharing a greater percentage of genetic material with each other than they do with each of their parents.

In Ackley and Littman's model, kin altruism can help to explain why the festival strategy is more successful in promoting altruism than the wind strategy. The continual reinvention of communication in cells occurs because of the high degree of relatedness of individuals in a cell; signals develop through kin altruism leading to a mutually beneficial relationship within a cell. After a windy day however, an unrelated individual enters the cell and, unbound by kinship, can exploit any evolved system to prosper at the other agents' expense. In contrast, after a festival day, a cell is unlikely to find itself

invaded by an exploitative agent. This is due to the nature of a festival, which will result in an agent being much more likely to mate if it comes from a population which is using the signals to communicate in a beneficial way, rather than exploiting others: cells are more likely to gain an agent who acts altruistically, as the cells where kin altruism is functioning are more likely to provide the mating agents. More generally, it can be seen that kin altruism can provide a useful explanation for the origins of coordinated communication systems where no benefit to the signaller seems to exist, however care must be taken to ensure that unrelated individuals do not have the chance to exploit the system.

Cangelosi & Parisi Cangelosi and Parisi (1998) present a model which they use to study the role of language as an aid to categorisation. At the core of their simulation is the idea that for a language to be useful to a population, the abilities to produce and comprehend signals must both evolve simultaneously. By invoking a scenario in which comprehension is rewarded while production is not, their model again addresses the issue of altruism. Specifically, they create a model in which an agent is placed in an environment with a range of mushrooms, some of which are edible while the rest are poisonous, and establish a system whereby agents receive signals from other individuals labelling the mushrooms, based on perceptual clues that the agent may be too distant to observe directly.

The simulation consists of a population of 100 agents, each of which is placed into an environment in turn. Within this environment, there are mushrooms which are either edible or poisonous, and can be discriminated between based on visual clues; these visual characteristics can only be observed when an agent is adjacent to the mushroom. Agents receive a fitness which is increased by the consumption of edible mushrooms and decreased by the consumption of poisonous ones. Neural networks form the basis of each agent, and map a set of inputs consisting of the location of the nearest mushroom and its perceptual properties (if it is sufficiently close to be observed), via a layer of hidden units, to an output that dictates the agent's next movement. Additionally, three input and output units are made available for communicative information to be

received and transmitted by respectively. Three strategies are examined for the use of these communication units. In the first, the input units are held at fixed values. In the second, an external language is provided, and the input units are set to given activation patterns depending on the type of the nearest mushroom, even when the mushroom is too far away to observe directly. In the final, a randomly chosen individual from the population is exposed to the perceptual properties of the mushroom nearest to the currently active individual, and its communication output units are used as input for the active agent's signal input units. The outcome of these strategies (which correspond to the situation in which there is no language, a fixed language and an evolving language) show that, after a sufficient number of generations, the population in which no communication exists is outperformed by populations where communication does exist. More interestingly, the populations with fixed language perform no better in the longer term than those with an evolving communication system, who quickly converge on a common pair of signals.

This work suggests a novel reason to explain the production of signals in the absence of benefit to the sender. Cangelosi and Parisi propose that the existence of a coordinated signalling system may be dependent on the architecture of the agents. In order for the received signals to be of use, they must produce distinctive activation patterns in the hidden layer of the neural networks for edible and poisonous mushrooms. Intuitively, it can be seen that evolutionary pressure exists to use one activation pattern for all edible mushrooms and another for all poisonous ones. Consistent activation patterns at the hidden layer will tend to produce consistent signals for the signalling units in the output layer: the coordinated output language develops as a result of consistent activation patterns at the hidden layer, which is a result of evolutionary pressure to correctly utilise input signals.

Noble *Costly signalling* is a further theory put forward by Zahavi (1975) to explain the origins of communication in situations involving a conflict of interest. Unlike the methods explored above to explain the evolution of altruistic communication, costly signalling is concerned with the evolution of honest signalling in situations in which

a signaller may seek to exploit a listener by lying. Noble (1999) presents computer simulations which seek to test the central claim of this theory, that honest signalling can be maintained in these situations by requiring a price to be paid in producing a signal.

Krebs and Dawkins (1984) suggest that in circumstances where exploitative signalling is possible, propaganda and advertising form accurate metaphors for the aim of the communicator. For example, in many species it is in the interests of the male to pretend to be a more attractive mate than it actually is; rather than accurately communicate its status to a potential mate, it attempts to advertise itself as superior to its real condition. Given a population of individuals who wish to misrepresent their knowledge in this way in order to gain some advantage, it might intuitively be expected that a communication system would eventually be abandoned due to the likelihood of receiving misleading information. Again, observation of the natural world reveals that this is not the case. Krebs and Dawkins suggest that the result will instead be an evolutionary arms race, in which receivers become increasingly hard to convince, and signallers must spend more time, energy and effort to be believed. This they call ‘expensive hype’ (and contrast it with ‘conspiratorial whispers’ that occur when communication is mutually beneficial). Zahavi’s theory of costly signalling is in agreement with this assessment.

Noble (1999) tests these predictions using a simple signalling game played by pairs of agents selected from a population of genetically evolving agents. The game consists of first randomly determining the state of the environment as either ‘high’ or ‘low’, allowing the signalling agent to choose whether to produce no signal, or a signal with a cost up to C_S based on the environment state, and allowing the receiver to respond if the signal cost exceeded its threshold T , at cost C_R . If the environment state was ‘high’ and the receiver responded, signaller and receiver receive pay-offs of P_S and P_R respectively, regardless of the cost of any signal sent. During each generation, 500 games were played using randomly selected pairs of agents, with an agent’s fitness being determined by its performance in all games it was involved in. The results obtained uphold the predictions made by theory; the threshold T decreases as the parameter P_R is set to increasing values, as does the cost of the signal produced. Signalling behaviour

only evolves in situations in which communication is mutually beneficial, however in situations in which a correct response by the receiver results in a large benefit to the signaller but no benefit to the receiver, the most costly signs are produced in order to attempt exploitation.

4.1.2 Compositional Communication

Hashimoto & Ikegami Few attempts have been made to model genetically encoded communication systems incorporating grammar, perhaps because such systems are difficult to find in the natural world. While lexical communication systems of this type correspond strongly to animal communication systems, and learned systems model simplified human language, few, if any, natural communication systems incorporate innate syntactic rules. Nevertheless, Hashimoto and Ikegami (1996) construct a model to study such a system. This model consists of a small population of agents, each of which contains a grammar, initially allowing the recognition of only a single terminal. At each timestep, each agent produces a set of 10 sentences, if able, which are given to the other agents to attempt to recognise. Agents' fitness is increased if they produce words not recently heard and for recognising the words produced by others using few rewriting steps, and the fitness is impaired by having ones utterances recognised by others. Mutation takes place during reproduction by adding or removing a rule from the grammar, or by adding, modifying or removing terminals or non-terminals from one of the agent's rules. The mutation scheme restricts the grammar to a context-free level at its most complex. The results of the model show that the proportion of sentences that can be recognised grows with time, in a series of 'punctuated equilibria'.¹⁴ However, the agents never reach the maximum expressiveness, as the ability to quickly parse and generate a small set of words common to the majority of agents becomes an overriding factor acting against the drive to greater generality.

The shaping of the population by the communication system, at the same time that the

¹⁴By limiting the length of generated strings, the language is kept to a finite size, hence this proportion can be calculated.

system is shaped by the population, mirrors the findings of Werner and Dyer (1992), discussed above. As in this earlier model, fitness is reliant on the behaviour of the rest of the population, as well as the individual's behaviour. In fact, the punctuated equilibria can probably best be explained by the population being unable to progress until enough agents fortuitously gain a similar, more expressive grammar through mutation at the same time; unless several agents mutate in a similar way at the same time, more expressive agents will be disadvantaged relative to the general population as they will have difficulty in quickly comprehending words commonly used by the rest of the population. Hashimoto and Ikegami draw analogies between the set of common utterances and human language. In particular, they remark on the similarity of the need to recognise members of this set quickly, and the comparable need to do so with common words in language. For example, common idioms are usually spoken as complete phrases, and may not actually make sense if their interpretation is attempted on a word by word basis. This model suggests that there are advantages in such phrases being stored in human grammar as a single rule. Recent work suggests that this may be the case (Wray, 2005).

Zuidema & Hogeweg The model of Hashimoto and Ikegami (1996) is modified by Zuidema and Hogeweg (2000) such that rather than the fitness function used in the initial model, they implement one of two scoring schemes, one of which rewards understanding others and being understood oneself, while the other rewards only understanding of others (thus implicitly penalising being understood). Additionally an 'innovation pressure' is sometimes added, rewarding novel utterances. The entire population is replaced at each timestep with children of agents from the previous generation, using the relative fitnesses to determine which agents will have offspring, as in most genetic algorithms.

When being understood is advantageous, it is found that expressive communication can be preserved or can develop from an initially empty grammar if innovation is rewarded, but will not develop from an initially lexicalised grammar. In contrast, when being understood is disadvantageous, expressive grammars quickly develop from all

initial conditions, even without the need for explicit innovation pressure. The increased expressiveness in the latter case likely develops from needing to understand what others produce, but create novel utterances oneself; the optimal way to perform this is to possess a language of which all other used languages are subsets. A notable aspect of the obtained results is the inability of expressive grammars to develop from an initially lexical language when communication is mutually beneficial. This is important as ‘*a large lexical proto-language with communication benefits for both speaker and hearer*’ is ‘*traditionally thought to precede the emergence of grammar*’ (Zuidema and Hogeweg, 2000). Clearly if this was a stage in the evolution of language, the pressure for grammar must have arisen from some factor outside the scope of this model, such as memory constraints.

Zuidema and Hogeweg further develop the model to look at altruism. If the fitness is calculated based only on the ability to comprehend others, speaking constitutes an altruistic act. Thus an evolving parameter is introduced into the model recording each agent’s probability of speaking, and as predicted by the theory of natural selection, it is found that this parameter rapidly tends to zero. However, when agents are placed in a 2-dimensional grid and are restricted to communicating with their immediate neighbours, this propensity to speak increases to nearly 100%, along with the expressiveness of the evolved language.

Altruistic communication in this situation can possibly be explained in part by recourse to kin altruism, as the model may make it more likely that closely related agents are closely placed, but is probably mainly due to *reciprocal altruism*. This explanation for altruism was developed by Trivers (1971) and in addition to explaining behaviour such as that shown here may be responsible for interspecies altruism where kin altruism clearly cannot be the cause. Reciprocal altruism may occur in situations in which repeated interaction happens between the same individuals. In such cases, while in a one-off encounter selfish action may be best, the repeated contact may mean a higher gain for both individuals involved if cooperation is engaged in. For example, some species of fish have developed relationships with ‘cleaners’ from other species. The cleaners are allowed to remove parasites from the fish, including in their mouths, while

		Player 1	
		Cooperate	Defect
Player 2	Cooperate	(3,3)	(0,5)
	Defect	(5,0)	(1,1)

Figure 8: Pay-off matrix for the prisoner's dilemma game.

the host fish reciprocate by allowing this without attempting to eat the cleaners.

The classic example of reciprocal altruism is in the *iterated prisoner's dilemma*, which Axelrod (1984) among others has studied. In the *prisoner's dilemma*, two individuals independently choose between two possible actions, 'cooperate' or 'defect', receiving a pay-off as shown in figure 8. Examination of the pay-offs reveals that whatever course of action is taken by one agent, the other agent will receive a higher pay-off by defecting than cooperating. The rational outcome therefore is for both agents to defect. In the related iterated prisoner's dilemma, an unspecified number, n , rounds of the prisoner's dilemma are played.¹⁵ In this case, if the strategy from the prisoner's dilemma is followed, both agents will obtain a final score of n , however, if both agents continually cooperate, they will both receive a greater pay-off, which is the maximum possible combined pay-off. In order for this strategy to be possible, both agents must engage in reciprocal altruism; in each round agents could gain a higher reward if they defected, but in the game as a whole they gain if they are prepared to act altruistically. If one agent attempts to 'cheat' by defecting at any stage, the other agent is able to punish this behaviour, by continually defecting from that point onwards; while the cheater will still score more highly than the other agent, it will score less than it would have done by acting altruistically throughout.

The fact that reciprocal altruism is rewarded in the iterated prisoner's dilemma is due to the relative values of the pay-offs, and clearly not all systems involving repeated interaction result in reciprocal altruism emerging. However, in Zuidema and Hogeweg's

¹⁵It is essential that the number of rounds of the game to be played is not able to be determined by the individuals involved, otherwise they could alter their strategy to exploit this fact.

model, the method used to calculate fitness will reward agents who act altruistically to their neighbours. If their environment consists of two spatial clusters of agents, one with the desire to produce utterances and the other remaining silent, the communicating agents will receive higher fitnesses than the non-communicating agents, because their neighbours produce output for them to observe. This leads to the number of communicating agents increasing until they come to dominate the environment. Similarly, should a mutant arise in the midst of a cluster of communicating agents, it will initially do well, but in doing so will create a small cluster of descendents who do less well than their communicating neighbours, who have more communicating agents adjacent. In contrast, when no spatial configuration is used, agents who communicate are no more likely to interact with other communicators than agents who remain silent and therefore reciprocal altruism cannot be established.

Reciprocal altruism offers yet another means of explaining apparently altruistic communication acts as actually being non-altruistic when viewed from another perspective. By acting altruistically in the short-term, individuals may actually be acting in their own best interests in the long-term. For example, in the model of Zuidema and Hogeweg, though agents apparently act altruistically in producing signals, the spatial configuration means any clustering of signalling agents are actually operating in a mutually beneficial way to increase their own fitness relative to groups of silent agents. Analogous benefits may help to explain the emergence of language or animal communication in the real world.

4.2 Cultural Evolution - Iterated Learning

The *iterated learning model* (ILM) is the single most widely used technique in studying culturally-acquired languages, and perhaps the most utilised method of all in modelling language evolution. Though earlier models (such as Oliphant and Batali (1997)) used similar techniques in learning lexical languages, Kirby (1998) was the first to employ it to study grammatical language, and it is with him that this technique is usually identified. Kirby has further developed the use of this model (Kirby, 2000, 2002a), as

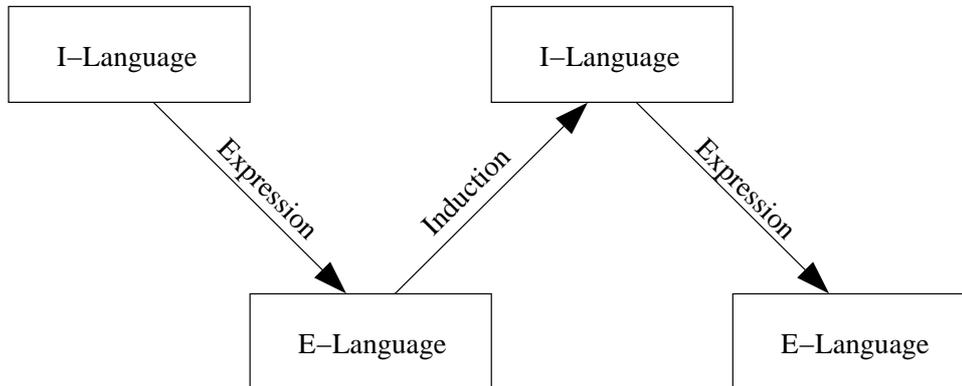


Figure 9: The Expression/Induction cycle. Language expressed by one individual is used to inform another’s knowledge of language through induction.

have others, including Brighton (2003) (who demonstrates the role of MDL in language learning), Smith (2003b) (who extends it to study the interplay between evolution and learning), Smith (2005) (who applies it to meaning creation) and Hashimoto and Nakatsuka (2006) (who explore grammaticalisation).

At the heart of the ILM is the idea of the *Expression/Induction cycle* (E/I cycle) (Hurford, 2002), which is closely related to the Chomskyan concepts of E- and I-language (Chomsky, 1980). As shown in figure 9, the I-language of one individual is used to express a set of utterances which are then used to modify the I-language of another individual through induction. While this cycle may occur in a group of peers, with each modifying their I-language in response to the others’ E-language, this model also fits the case where transmission is vertical, from one generation to the next. It is behaviour in this cross-generational situation that the ILM seeks to capture.

Population turnover is as essential in an ILM as it is in studying genetically encoded communication. While individuals in an iterated learning model acquire an internal knowledge of language through learning based on the expressions created by other, older agents, this learning typically takes place only once when the agent is created. Due to this, it is necessary to continually remove older agents from the population and replace them with newer ones, as the inexact induction from the population’s current language is the process that drives language change. Two forms of ILM exist, which

refer to the method by which the next population of agents is formed, the *generational* model and the *gradual* model. In the former, a completely new generation is formed at each timestep, inducing its internal languages from the expressions produced by the previous generation. In the latter, a single agent leaves the population at each timestep, to be replaced by a new individual, who acquires language through induction of the expressions produced by the remaining population. Removal is typically through random selection or chronological ordering, with no regard given to linguistic competence.

While the ILM fails to capture the full complexity of human linguistic learning, it offers a useful abstraction for theories relating to language evolution. Despite the ability to continually learn new words and even languages throughout life, the majority of linguistic learning, particularly of grammar, occurs during youth: language learning in humans functions similarly to an ILM, in which the bulk of learning occurs when an agent is added to the population, with subsequent minor alterations being permitted. Accepting that most models neglect the later modification of linguistic competence, the ILM offers a framework sufficiently close to the method whereby human language is learned to provide useful insights into the processes that may have driven the evolution of language.

4.2.1 Lexical Communication

Oliphant & Batali Drawing inspiration from the alarm calls of the vervet monkey, Oliphant and Batali (1997) produce an ILM in which gradual turnover of a population leads to coordinated signals emerging to unambiguously express particular meanings. The vervet monkey (Cheney and Seyfarth, 1990) is prey for three species of predators and very different action is needed to evade each one. Vervets have evolved to produce distinctive alarm calls for each danger, alerting the rest of the group to the presence of a predator and allowing the correct action to be taken. While in the case of vervets, this signalling appears to be innate, Oliphant and Batali use this example as the starting point from which to explore the question of coordinating learned signalling and reception behaviour in a population from an initially random situation. In the model that

they produce, they do not consider the altruistic aspect of producing alarm calls, an issue perhaps better studied in a genetically evolving population.

The linguistic competence of each individual in the model is determined by two functions, a *send* function and a *receive* function, probabilistically determining which signal is transmitted for which meaning, and what meaning should be understood for each signal received respectively. These functions are not necessarily consistent: an individual may most often emit a signal in a given situation, but may associate this signal with a different meaning on hearing it. The model evaluates the performance of several learning algorithms for initialising an agent's functions when it is initially added to the population.

The first algorithm considered sets each new agent's functions to be those of the population average (which can be approximated by regarding an arbitrarily large number of communicative episodes). This produces populations capable of maintaining an optimal communication system but not of improving a sub-optimal system. By modifying this learning method such that new agents acquire a *deterministic* approximation to the population average, the population is able to acquire a communication system with a better communicative accuracy given a reasonably coordinated system, but if the language is initially random, it may degrade even further due to the accumulation of homonyms and synonyms. Finally, the *Obverter* learning procedure is considered. Under this scheme, rather than acquiring a system that replicates the average communicative behaviour of the population, an agent learns so as to maximise its communicative success when interacting with the population. This is done through setting the receive function based on the population's send function, and the send function based on the population's receive function. As with the previous algorithm these learned functions are made deterministic. Not only does the use of the Obverter function improve the communicative accuracy to a given maximum in practice, but Oliphant and Batali are able to mathematically prove that it must do so theoretically too. Approximations to the Obverter algorithm, in which the more realistic situation of only allowing a small finite number of observations to be used to estimate the average population behaviour, also result in high quality communication systems developing, though with declining

performance as the number of observations decreases as would be intuitively expected.

This model demonstrates the importance of the algorithm used to induce a new agent's knowledge of language in an ILM. Despite the fact that the same expressed behaviour was available in all experiments,¹⁶ and that population turnover occurred in the same way, the emergent behaviour exhibited by the system varied greatly depending on the learning algorithm used. This has clear implications for modelling human language with an ILM. While one could conceivably obtain realistic data for the expressions used to set the I-language, the resulting output could not be accurately modelled without knowledge of the learning algorithm used to process this data.

Smith As with the overwhelming majority of studies using the ILM, the individuals in Oliphant and Batali's model have access to both the signals and the meaning they were used to convey in the data from which they learn. However in real life, learners usually only receive the signals and must infer the intended meaning from the context in which the communication occurred. Working from the insight that if the meanings are transferable, then establishing coordinated signalling is redundant, Smith (2005) uses an ILM to investigate how coordinated communication can be established in a more realistic setting, in which individuals may only determine meaning from context.

Smith places a model of inferential learning that he has previously developed (Smith, 2003a) into an ILM, with the goal of monitoring not only the transmission of the lexicon from one generation to the next, but also the semantic concepts that the lexicon is able to express. In this model, agents are not given initial meanings, but create them by attempting to disambiguate between objects in the world. Each object produces perceptual properties (consisting of a real-valued number between 0 and 1) on a number of channels, and agents create meanings that consist of intervals within these channels. To distinguish an object from others, an agent must identify one of its own existing meanings such that the object's property falls into the given range but those of the remaining objects do not. If this fails, an agent will split an existing meaning into two

¹⁶Except, of course, for those experiments in which the number of observed episodes was deliberately restricted, to assess the role that this played.

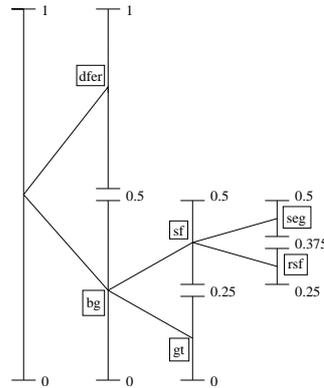


Figure 10: A meaning tree for a single channel, as used by Smith (2005). Specificity of meaning increases with depth. Numbers show the range of values for each meaning, and associated words are shown as boxed text.

for future use, such that the required meaning can now be expressed. In this way, a tree structure of distinguishable meanings is built (see figure 10). This representation of meaning is based on that developed by Steels (1996a).

Agents in the ILM undergo two processes of learning when created instead of the conventional one: an initial meaning creation phase followed by the traditional language learning one. In the initial process, they are repeatedly exposed to objects in the environment which they learn to distinguish from a context formed of other objects by the method of meaning creation outlined. Following this they are taught language by observing a number of communicative episodes by the agent in the previous generation. In contrast to most other models however, they are not told the meaning or object to which the given word refers, only the context in which it was used; a process of inference is used to associate words with meanings based on their possible meaning in the given context.

Results are obtained for the stability of the lexicon through the generations. While communicative accuracy remains high between generations, the persistence of individual words reaches only around 80% between consecutive generations and the proportion of the original language remaining falls even faster, to around 20% after only 10 generations. Two processes appear to cause this. Firstly, a ceiling on the number of words which can be learned is formed by the meaning formation stage: agents can only

later infer that a word pertains to a particular meaning if they originally learned that meaning. Secondly, the number of words correctly learned is further limited by errors committed during the inferential stage.

In seeking to explain these results, Smith inspects the agents' lexicons. This reveals that the words which persist across generations are those which refer to the most general meanings. Two reasons for this can be identified. The tree structure of the meaning space results in most agents acquiring the most general meanings of each channel, that is, those meanings that most coarsely divide the channel. This makes it possible for them to later correctly learn the words for these meanings. This contrasts with very specific meanings which are likely to be created only if an agent is exposed to a particular small number of examples requiring their use. Secondly, as agents use the most general meaning possible for discrimination in the language expression phase, learners are likely to observe more examples with these general categories than with very specific ones. To reiterate, the words associated with the most general meanings are more persistent because these meanings are both more likely to be necessary and are used more often. This creates a prediction that could be tested in human language; should these findings be more than just artefacts of the model, it would be expected that similar results could be found in real languages. Such evidence may exist in English, which despite having incorporated loan words from many languages, most extensively those descended from Latin, retains Germanic vocabulary for its most commonly used words.

4.2.2 Compositional Communication

Kirby It is through the work of Kirby (2002a) that perhaps the most significant use of the ILM has been made. An ILM is created in which generations, consisting of a single agent, induce then express signals (resorting to invention if necessary) to describe a subset of the possible meanings drawn from a semantic space, expressed in first-order or higher logic. The induction algorithm used by the individuals is highly complex, especially with respect to rules for preventing over-generalization, but operates as to associate parts of a signal found in multiple examples with the common semantics from

those examples. The results produced after several thousand generations of this model are very interesting; a fully compositional language develops eventually in all situations. Even more strikingly, the mechanism that causes the development of compositionally has conventionally been seen as a problem in learning language. To understand the procedure by which a compositional language emerges from this model, it is necessary to examine the development of the grammars of the agents over time.

Early in the simulation, agents possess grammars that map whole phrases to strings, such as that shown in figure 11. This grammar lacks any structure: the two *detests* predicates (rules 1 and 5) have no commonality in their representation and common occurrences of objects are not reflected in common occurrences of substrings.

Later in the simulation, agents develop a language that is partially compositional, such as the grammar of figure 12, which shows some signs of regularity. For example when the second argument is *mary* (rule 8) then the predicate and first argument have become lexicalised, though the grammar is far from entirely compositional.

Finally, the simulation stabilises at a language that is fully compositional and consists of the minimal number of rules, with each atomic item appearing only once in a lexical rule. However, despite the grammar consisting of the minimal number of rules, the rules are more complex than necessary: they contain extra characters such as the *v*, *g* and *n* in rule 13.¹⁷

The underlying mechanism that causes this increase in compositionality is the so-called *transmission bottleneck* that prevents an agent observing signals for all meanings in the semantic space when learning. This bottleneck prevents languages lacking compositionality from accurately being transmitted from one generation to the next: some

¹⁷The extra characters present in the rules are an artefact produced by details of the simulation. In a very similar study, Hurford (2000) investigates the problem of language acquisition, modifying the model to include a population of agents in each generation and employing a synthetic process, rather than an analytical one, to learn rules: agents in this system build complex meanings from simple meanings, not induce simple meanings from complex as in Kirby's model. Hurford's model does not evolve semantically empty strings in the compositional rules unlike Kirby's, which are caused by the fortuitous effects of the induction process.

$$\begin{aligned} \mathbf{S}/\textit{detests}(\textit{john}, \textit{gavin}) &\rightarrow \textit{nqb} & (1) \\ \mathbf{S}/\textit{hates}(\textit{heather}, \textit{mary}) &\rightarrow \textit{b} & (2) \\ \mathbf{S}/\textit{loves}(\textit{mary}, \textit{pete}) &\rightarrow \textit{k} & (3) \\ \mathbf{S}/\textit{admires}(\textit{john}, \textit{mary}) &\rightarrow \textit{u} & (4) \\ \mathbf{S}/\textit{detests}(\textit{pete}, \textit{john}) &\rightarrow \textit{ayj} & (5) \\ &\vdots & \end{aligned}$$

Figure 11: Example of a grammar that may be found near the beginning of an experiment by Kirby (2002a). Non-terminals are shown in bold, with italicised terms on the left-hand side corresponding to meanings and italicised terms on the right to (parts of) words.

$$\begin{aligned} \mathbf{S}/\textit{hates}(\textit{pete}, \textit{john}) &\rightarrow \textit{a} & (6) \\ \mathbf{S}/\mathbf{p}(\textit{john}, \mathbf{x}) &\rightarrow \mathbf{A}/\mathbf{x} \mathbf{B}/\mathbf{p} & (7) \\ \mathbf{S}/\mathbf{p}(\mathbf{x}, \textit{mary}) &\rightarrow \textit{l} \mathbf{B}/\mathbf{p} \mathbf{A}/\mathbf{x} & (8) \\ &\vdots & \\ \mathbf{A}/\textit{gavin} &\rightarrow \textit{b} & (9) \\ \mathbf{A}/\textit{mary} &\rightarrow \textit{ni} & (10) \\ &\vdots & \\ \mathbf{B}/\textit{loves} &\rightarrow \textit{y} & (11) \\ \mathbf{B}/\textit{hates} &\rightarrow \textit{n} & (12) \\ &\vdots & \end{aligned}$$

Figure 12: Example of a grammar that may be found half-way through an experiment by Kirby (2002a). See figure 11 for a description of the notation used.

meanings will need to be expressed that have no associated signal (leading to invention) and some meanings that are paired with signals will not be used (leading to loss). This leads to changes in the language in every generation. From these initial, unstable languages, compositionality enters the language as occasionally, due to chance, a learner will be given several very similar meanings with similar words associated with them, due to random inventions by the speaker. When this occurs, the listener will construct a rule which is used in the production of several different meanings, and therefore will have a greater probability of being used during this agent’s subsequent expression phase, in turn increasing its chance of being present in the language of the

$$\mathbf{S/p(x,y)} \rightarrow v \mathbf{A/y} g \mathbf{A/x} \mathbf{B/p} n \quad (13)$$

$$\mathbf{A/gavin} \rightarrow gw \quad (14)$$

$$\mathbf{A/john} \rightarrow gbb \quad (15)$$

$$\mathbf{A/pete} \rightarrow k \quad (16)$$

$$\mathbf{A/heather} \rightarrow gyt \quad (17)$$

$$\mathbf{A/mary} \rightarrow pd \quad (18)$$

$$\mathbf{B/hates} \rightarrow a \quad (19)$$

$$\mathbf{B/loves} \rightarrow c \quad (20)$$

$$\mathbf{B/detests} \rightarrow m \quad (21)$$

$$\mathbf{B/admires} \rightarrow srw \quad (22)$$

$$\mathbf{B/likes} \rightarrow z \quad (23)$$

Figure 13: Example of a grammar that may be found at the end of an experiment by Kirby (2002a). See figure 11 for a description of the notation used.

next generation. Over time, those rules which are involved in the derivation of the greatest number of meanings will be more regularly used, and thus more likely to be re-induced by the next agent. Eventually this process leads to the establishment of a stable, compositional language with the minimum number of rules.

The most remarkable aspect of this research is that it converts a potential problem that exists in language acquisition into the driving mechanism behind the evolution of compositionality. The transmission bottleneck which leads to composition emerging in the language is actually part of the poverty of stimulus argument (as summarised in Pullum and Scholz (2002)) which argues in favour of language universals being largely innately defined. Kirby's work suggests that the languages existing in the world today have evolved to be those that can be learned from a very small number of examples, in contrast to the nativist position that it is humans who have evolved to learn languages more easily: the findings reverse the problem and show that the need for languages to be learned from few examples creates selective pressure on languages to evolve, and not necessarily on humans.

Smith (2003b) further investigates the foundations of the ILM and generalises Kirby's work to show that language can only reliably emerge from an ILM given three factors;

- A transmission bottleneck must exist
- The environment possesses some regularity or structure
- The algorithm used for learning must be biased
 - in favour of one-to-one mappings
 - to utilise regularities in the environment

The relatively weak constraints that these findings place on the learning algorithm provide further evidence that language acquisition may not require a specialised language learning device in the brain, or at least, that early language users did not need such a device.

4.3 Cultural Evolution - Negotiation

As with the ILM, the negotiation model seeks to simulate the evolution of communication as a cultural process in which communicative behaviour is acquired based on the behaviours of others, in line with the E/I cycle. However, in contrast to the ILM, which seeks to explain the evolution of language through the vertical transmission of language from one generation to the next, the negotiation model investigates the establishment of a coordinated communication system within a fixed group of agents,¹⁸ who continually adjust their behaviour based on that of the others. Whereas agents in an ILM engage in a period of learning when entering a population, and retain their competence unmodified thereafter, in the negotiation model, agents enter the initial population with no linguistic knowledge, and continually add to and revise this knowledge based on their success in communicating with the rest of the population.

Both the ILM and negotiation are capable of potentially providing insight into the processes behind the emergence and evolution of language, though undoubtedly, the

¹⁸In fact, it is possible to extend the model to allow for the entrance and egress of agents during the experiments. However, whereas in the ILM this is an indispensable part of the process, in the negotiation model it disrupts the coordination achieved in the language and is implemented, for example, in order to establish how robust the language is to the disturbances caused by population change (Steels, 1996b).

ILM has become the dominant paradigm. The ILM shows how language may change through the generations as a result of the way in which it is transmitted through the E/I cycle, but language change can also happen at a much faster rate; in the modern world, words such as slang or technical expressions are continually entering languages and spread over a timescale measured in months or years, not generations. The ILM is incapable of modelling language change on this time-scale, but negotiated models are able to address this issue. Furthermore, the ability to produce a coordinated language within a single generation, as in the negotiated model, questions the established position that language evolved in stages over the course of many generations: results obtained with the negotiated model show that, should the correct cognitive abilities be in place, full language may arise in a population within the limits of a single generation, a theory supported by real world evidence (Senghas and Coppola, 2001).

In building a negotiated model, it would be trivial to construct a setting in which a population of agents arrived at the same communication system after all observing the same small set of linguistic interactions. However this forms a poor model of the real world. In reality, each individual is only ever party to a minuscule proportion of all the acts of communication that have occurred. Despite this, groups of millions of individuals exist who are able to conduct a conversation using language sufficiently similar to each other to be mutually intelligible. It is the major goal of studies involving negotiated models to investigate how the communicative activity may be coordinated at a universal level by modifications to behaviour based solely upon episodes involving subgroups of individuals. In practice, this is carried out through agents repeatedly engaging in *language games*, small communicative episodes in which a pair of agents interact within a well-defined conversational framework and modify their future linguistic output based on the success of the communication.

4.3.1 Lexical Communication

Steels Numerous issues have been explored by Steels through the use of negotiated models, including dialect formation (Steels, 1997), embodiment (Steels and Vogt,

1997), concept formation (Steels and Kaplan, 1999) and spatial vocabulary development (Steels, 1995). In all cases, a core model based on a particular language game structure is used, modified to focus on the object of study.

A successful language game may proceed as follows.

- An agent (called the initiator) chooses an object of discourse from those available.
- The initiator forms a discrimination set that unambiguously indicates the chosen object, and lexicalises each property in that set using its vocabulary.
- A second individual (the responder) interprets the words, and chooses the object to which they refer.
- The responder articulates a discriminating description of the object back to the initiator.
- The initiator transmits an extra-linguistic signal indicating that the received description matched the object initially described.
- Both agents increase the weighting that they give to the association between the words used and the meanings they were used to represent.

While successful dialogues such as the one above do occur, most of the above steps contain the potential for the conversation to fail in at least one way. This usually leads to one of two actions; it is possible for new associations to be formed between words and meanings (for example, if a meaning is to be expressed but no word is known for it), and associative weightings between existing words and meanings may be reduced (if a word set was misinterpreted). Initially most games fail (due to the initially blank lexicons possessed by agents), but these failures lead to important developments by permitting the creation and diffusion of words associated with particular meanings. As the number of communicative episodes increases, the communicative accuracy of the population increases as agents converge onto a single word to express each meaning, and ambiguity in the meaning of words is resolved. Eventually, complete coherence of language is observed.

In some settings, the effect of introducing additional objects or agents into a system in which the existing agents had converged onto a common lexicon was assessed (Steels, 1997). The introduction of additional objects (even in large quantities) proved to have little effect on communicative accuracy, reducing it temporarily while new words were invented and circulated. Likewise, the entrance of new agents led to a drop in communicative accuracy for a relatively short period, until they acquired the language already dominant in the population. Removing agents from a population proved to have no effect on the communication between the remaining agents.

The effect of spatial distribution of agents on the language they develop has also been assessed (Steels, 1995). In this study, agents were arranged on a two-dimensional plane, with a probability of conversing inversely proportional to the distance between them. By placing agents so as to form three distinct clusters, a model of communities with individuals conversing significantly more with those inhabiting the same cluster than outsiders was formed. Within each cluster, coherence reached high levels, but between clusters remained relatively low. This confirms expectations arising from real life: in situations in which continual infrequent contact between individuals from different communities occurs, each language may admit an occasional word, but do not converge into a single language. The effects of increased contact between speakers of different languages could also be modelled, by decreasing the effect of distance after intracluster coherence was reached. This led to the rapid development of bilingualism, followed by the mixing of the languages towards a single language formed by incorporating words from each of the three original languages. Again, these results correspond to real world observations (Bickerton, 1975).

4.3.2 Compositional Communication

Batali A negotiated model is seldom used for the task of studying grammatical language, probably in part due to the successes experienced in simulating this phenomenon using the ILM. One of the few models that does employ this technique in this role is that presented by Batali (2002), which successfully demonstrates convergence to a shared

grammatical language, albeit one with an unusually represented internal grammar.

The semantic space employed by Batali is slightly unusual, and is motivated by a desire to demonstrate recursive language without using agents with explicitly recursive internal representations, ‘*Any recursive regularities in the systems that agents develop to express these non-recursive meanings will be entirely of their own invention.*’ Therefore, the semantic space consists of two types of predicate. The simpler type is a unary predicate, such as *snake*(1), which states that individual number 1 has the property of being a snake. The other type are binary predicates, for example *bit*(1, 2), which states that individual 1 bit individual 2. Meanings can be produced by combining sets of these predicates to form complex relationships between multiple objects.

Batali’s model is based on a population of agents who, rather than induce a grammar from the communicative behaviour they observe, follow the lazy learning method of case-based learning (Aha et al., 1991), and simply store the behaviour as a collection of exemplars. Initially, learning will consist of each agent simply storing the meanings referred to in conversations of which it was part, along with the phrase used to describe it and a score that records the phrases usefulness in understanding new phrases. When agents gain more phrases, through chance some phrases will be observed in which a subset of the meaning has already been observed along with a subpart of the phrase (as in Kirby’s ILM (Kirby, 2002a)). At this stage the learner can acquire not just the new phrase, but also induce a method of constructing this phrase from the subphrase it had knowledge of and hence infer a meaning for the rest of the phrase. The subphrase is stored along with a trace of how the phrase was (de)constructed.

When called upon to express a meaning, an agent will create a phrase for the given meaning. If a meaning is already associated with a phrase in an exemplar, then this will be used, otherwise the meaning will be created through modification of existing exemplars as shown in figure 14. Modifying an exemplar can consist of either replacing a subset of the meaning (and hence a subphrase as well) with another exemplar, or combining two exemplars (resulting in the concatenation of their phrases and meanings). Both these options have associated costs, with the phrase with the lowest cost

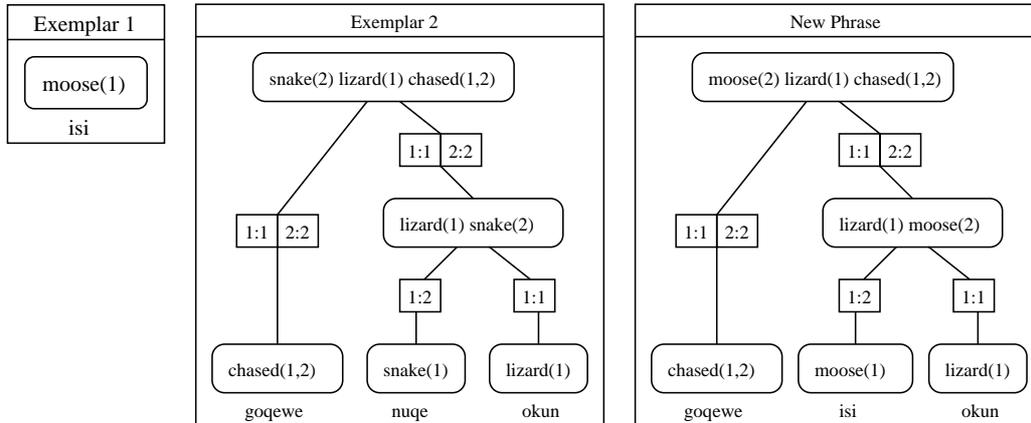


Figure 14: An example of combining two exemplars from Batali (2002) to produce a new phrase. The phrase expressed in the third box is created by replacing the part of the meaning and phrase from the exemplar in the second box with that of the exemplar in the first box.

of generation being chosen from all candidate phrases. If there is no way to form the required phrase out of an agent’s existing exemplars, a new exemplar for that meaning will be created.

Agents quickly accumulate a large number of exemplars, many of which are of little use in generating new sentences. This problem is solved by the twin methods of pruning and reinforcement. Pruning of exemplars is handled very simply. Any exemplar that has not been used by an agent for 200 episodes is removed. The paper notes that a useful exemplar tends to be used every 30 turns, making phrases which are pruned almost certainly unnecessary. Reinforcement is used when an agent ‘parses’ a phrase. If an exemplar is used in this process and the agent was able to recover the correct meaning from the phrase, then the exemplar receives positive enforcement and has its cost slightly reduced. An exemplar that was used in a failed parse has its cost slightly increased.

The unique nature of the learning algorithm and semantic representation make direct examination of the results complex, but some factors can be identified by comparing the results of this model and that of Kirby (2002a). In Kirby’s system the grammar initially is very large, before contracting to the smallest number of necessary rules. A similar process is true of this work too. Creation of new exemplars is initially very common

(analogous to the initial random creations of Kirby), but then the number of exemplars rapidly falls to leave a much reduced core set with almost negligible cost, indicating very common use. One clear distinction between the models is that Batali's results lack anything that would be traditionally recognised as a grammar, being merely a set of exemplars with costs. However, as Hurford (2002) notes, this is an almost trivial difference; all that is missing is a final induction step that transforms the 'grammar' of Batali's simulation into a traditionally expressed grammar. A more important difference is that, like Kirby, this model generates strings without direct semantic meaning, but, unlike in the other model, these have roles as modifiers or as delimiters. For instance, the less conventional structure of the semantic space leads to a situation in which strings without their own semantic meaning arise to stand for such functions as making a verb reflexive (i.e. mapping *bit(1,2)* to *bit(1,1)*). It is also possible to find similar strings that operate to switch the situation from a creature being the object of an action to being its subject. Though arising from similar random processes to the semantically empty strings in Kirby's work, these strings appear to be far more significant: the role that they play in the simulation seems to mirror that played by verb endings or case markers in human language.

4.4 Summary

This chapter has presented a structured overview of currently existing research involving the simulation of communication with computational models. Through the research presented, particular emphasis was given to models which have produced significant contributions to the understanding of language development and those which explored aspects of the evolution of language which are germane to the model developed in the remaining chapters of this thesis.

The issue of communication in situations in which the interests of the individuals involved conflicted was addressed, and two mechanisms to explain the existence of altruism in communication were outlined, kin selection and reciprocal altruism. The question of altruistic behaviour arises in later chapters due to the inherent, apparent

selflessness of the agents in the model, who utilise language to provide others with directions to resources. The model also assumes a pre-existing lexicon shared between individuals. It is envisaged that a series of language games, such as those developed by Steels (1996b) could be used one viable method used to create such a system.

5 Using Artificial Life to Model the Evolution of Language

The chapters of the thesis following this one will be concerned with experiments performed in order to assess the plausibility of a trait promoting syntactic communication being able to spread through a population through natural selection. In order to distinguish clearly between the specified behaviour deliberately built in to the model and the emergent behaviour that develops from it, the model itself will first be presented in this chapter. In describing the model, the principles will be outlined, motivation shown and theoretical analysis of the type of language concerned included.

A major component of this specification of the model is the introduction of the representation used by agents. In earlier chapters, it has been established that a common neurological structure may underlie syntactic and navigational abilities. The representation used for the experiments, and its duality in agents' navigational and language processing, is consistent with this and makes the link even more explicit.

The rest of this section proceeds as follows. First the computer model is described in broad terms. This is followed by explanations of the representation used by the agents for navigation and communication, and its utilisation for both tasks. The behaviour undertaken by agents in the model is then explained in section 5.3.¹⁹ As the implementation of the model runs to many thousand lines of code, all of this description take place at a relatively high level, allowing the concepts to be understood without obscuring them with details of the implementation. Following this, additional lower level aspects of the computational model used are given in section 5.4. The behaviour specified by the model is shown to be viable in section 5.5; an environment is presented in which a population of agents achieves a stable size through using knowledge to return to, or to communicate directions to, previously visited resources. The range of environment conditions in which this may be beneficial is explored in chapter 6. Finally, this chapter concludes with a theoretical discussion of the type of language permitted

¹⁹Variations of the basic behaviour will be outlined in chapters 6 and 7 as appropriate.

by the model, with specific emphasis on properties shared with human language.

5.1 The Computer Model

The purpose of this thesis is to present the notion that the ability to handle syntax in language may have been derived from earlier navigational abilities. While the survey of existing literature has shown that there are strong reasons to believe that a common neurological basis exists for these abilities, this only explains the origin in part. It is not sufficient to suggest that the neural circuits under-pinning syntax handling may have been exapted from those involved in navigation, by what may have been a relatively small mutation. This explains how a single creature gained the ability, but not why the mutation spread, through natural selection, to be found in every modern human; it explains the origins in an individual, but not in the species. It is crucial, therefore, to investigate how such a mutation could give the carrier a reproductive advantage over others without it. In order to do this, it is necessary to establish another theory to supplement the primary thesis that linguistic abilities derive from navigational ones.

Computer modelling operates by abstraction of the real world to create a simpler model which captures the important details, while leaving aside those which are irrelevant. Therefore, it is impossible to test whether the use of syntactic language offered any evolutionary benefit without knowing the purpose to which it was originally put: without some knowledge of what the language might be used to discuss, it is impossible to know which aspects of an individual, or the environment, to include in a model and which to abstract away. In light of this, the supplemental theory is proposed that the original purpose of language was for the communication of directions to resources needed for survival. The transient nature of language makes it impossible to ever definitively establish any such supposition, however in the present context (that of discussing the potential for linguistic abilities to be adapted from route-planning ones) it is not an unreasonable conjecture.

From this theory, it is relatively straight-forward to conceive an experimental setting

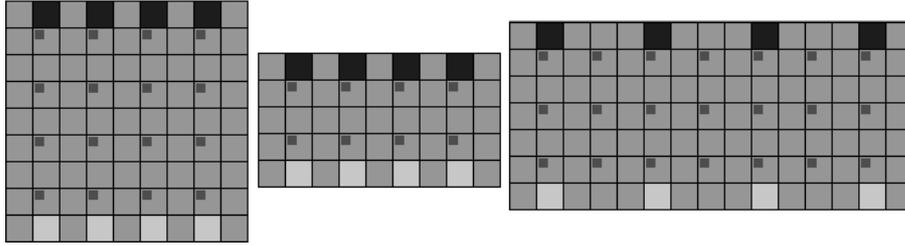


Figure 15: A range of environments with different spacing of resources. Each map contains an equal number of water resources (dark squares at top) and food resources (light squares at bottom), but with different distances between them. Landmarks are shown as regularly spaced boxes within the environment.

to test any reproductive advantages that may be associated with the use of syntactic language (Kazakov and Bartlett, 2002, 2004b). A population of agents are placed in a two-dimensional grid that forms a virtual world, in which there are also food and water *resources* which the agents must regularly utilise in order to survive. Agents collect from these resources when they find them and then consume their carried stores over subsequent turns, dying if they run out of either resource. *Landmarks* are also placed into the environment, relative to which navigational instructions may be given. Agents may move around the environment, searching for resources to survive and mating when meeting another agent and possessing sufficient stored resources to afford to do so. A range of environments for exploration can be created through positioning resources differently or by altering their volatility (see figure 15). Any benefits that syntactic language brings can thus be measured by any improvement in performance at surviving and reproducing by agents capable of communicating directions to resources syntactically over agents implementing different policies.

In choosing to assess the impact that communication has through the ability of agents to accomplish a simple task to aid survival, rather than evaluating communication complexity or convergence, the approach taken here follows the advice of Werner and Dyer (1992), who state;

We believe that a number of general principles should be followed when setting up simulations to evolve communication among organisms in artificial environments. [...] First, there should not be direct pressure on the ani-

mals to communicate. Communication should arise as a solution to another problem that has to be solved by the population. [...] Animals should not be judged on how well they communicate, but on how well they solve the task at hand. In this way, one can determine how communication aids in tasks that would normally be faced by an evolving population. [...] Second, it is important to present the populations with natural tasks such as finding food, protecting young, and attracting mates. These are the kinds of tasks that animals and humans faced when creating their communication systems, and such tasks placed important constraints on the development of communication.

The method used here adheres to both these principles. In the case of the first point, several possible behaviours are considered, with no guarantee that the most successful one will be the use of syntactic communication. Indeed, as will be shown in chapter 6, environmental factors such as the position or volatility of resources are hugely influential in determining the relative value of different strategies. The second point is trivially met through the design of the experimental setting.

There are actually two different experimental approaches that can be derived from the specification of the simulation described above. In the first, the environment can be inhabited in turn by a series of populations, each of which consists of agents who all implement a communicative behaviour common to that population, but with this behaviour varying between different populations. By running the simulations for enough time to allow multiple generations of agents to be created while recording the population sizes, it is possible to establish which policies offer advantages by comparison of the relative magnitudes of the populations. This is the approach taken in this thesis. The other approach, which is not investigated, is to initiate the simulation with agents possessing different communicative abilities present in a single population. In this case, the relative merits of a strategy can be observed through the proportion of the population exhibiting each behaviour in subsequent generations.

There are merits and drawbacks to both of these approaches. In the case of the first

approach, evaluation of the relative advantages of a particular strategy is simple. However, the fact that the strategy is only evaluated at the population level is a problem. While a particular policy may be best for the species as a whole, it may not be best for any given individual. Effectively, the approach is based on the largely discredited evolutionary idea of *group selection* (which states that evolutionary forces act at a species or group level) rather than the more modern neo-Darwinian framework. In contrast, the other approach deals directly with this problem by having heterogeneous behaviour in a single population; if a particular strategy is better than another, then those implementing the better one will have more offspring than those implementing the other, hence the proportion of agents following the better behaviour will rise over time. A problem with this approach however is that, as is the case with any social behaviour, the value of any strategy is dependent on the strategies of the rest of the population. This would make analysis of the results much more complex than those obtained through the first approach. It is due to this difficulty in evaluating and analysing the results that the former approach has been adopted.

5.2 Knowledge Representation

5.2.1 Aboriginal Songlines

The representation underlying the form of navigation and communication used in the experiments draws inspiration from *Songlines*, a form of oral tradition among the Australian Aboriginal tribes. Songlines reflect a belief that ‘*Ancestral Beings roamed once the face of the Earth, creating the features of the landscape [...] along the tracks that mark their journeys.*’ In the grouping of songs into series, ‘*the most pervasive is the geographical organisation of songlines, where each “small song” equates with a different site along an ancestral journey, and the series as a whole constitutes a sung map*’ (Barwick and Marrett, 2003).

It is important to note that, while the approach is inspired by songlines, the claim is not advanced that language evolved from this phenomenon. It is not a necessary part of

the argument that songlines represent an intermediate stage of language evolution, or indeed that they were ever more widely used by the human population outside Australia. They are presented here merely to show that the form of navigation proposed, and its link to language, is not entirely unrealistic, and has an obvious analogue in the real world.

5.2.2 Navigation

In creating the system of navigation used by agents, some of the features of the songline system are simplified and other aspects ignored completely. The result is a method of navigation in which individuals can travel between locations by remembering an ordered list of which landmarks it is necessary for one to pass in going from the initial location to the destination.²⁰ This representation is an extremely impoverished description of a route, containing the minimum amount of information necessary to correctly record a path and ignoring superfluous data, such as how to travel from one landmark to the next, which can be found through observation in the environment. More formally, the navigational information can be stored in rules of the following form:

$$\textit{goto}(\textit{food}) \Rightarrow \textit{goto}(\textit{PosX}) \quad (24)$$

$$\textit{goto}(\textit{PosX}) \Rightarrow \textit{goto}(\textit{PosY}) L_1 L_2 L_3 \quad (25)$$

$$\textit{goto}(\textit{PosY}) \Rightarrow \textit{goto}(\textit{PosX}) L_3 L_2 L_1 \quad (26)$$

where L_n are landmark descriptions and \textit{PosX} and \textit{PosY} are locations in the environment. The landmark descriptions are vectors of the exact physical properties of the landmark (allowing accurate recognition on subsequent reexposure),²¹ while the locations consist of an unordered set of the descriptions of all landmarks visible from that position, i.e. those within a given distance of the position.

²⁰There is evidence that some species navigate in a similar way; For example, wood ants ‘*tend to treat discrete landmarks along the route as beacons at which they aim (Nicholson et al. (1999)), making each landmark an intermediate goal*’ (Graham et al., 2003).

²¹Landmarks are chosen such that descriptions are unique. This means that each landmark’s description acts as a *de facto* identifier for that landmark.

The rules given above can be seen as defining equivalent actions, for example, rule 24 can be understood as stating that to go to a food resource it is sufficient to go to *PosX*. Rule 25 states that to go to *PosX*, one can go to *PosY* and then pass the landmarks described by $L_1 \dots L_3$ in order, and rule 26 is an inverse of the second, allowing this path to be traversed in the opposite direction.

These rules are initially accumulated by agents through exploration of the environment. While roaming the environment, each agent collects and stores information about the landmarks it passes (as well as the last location for which it stored a rule) in a short-term memory that is distinct from the set of rules that it maintains. When a position of some importance is reached, this sequence of landmarks is used to add a new rule to the agent's knowledge base stating the path between the locations in question. In fact, the above set of rules are those that would be formed by travelling from *PosY* to *PosX* and discovering food at *PosX*; rule 24 is added to record why *PosX* is significant and rules 25 and 26 show how to get to this location from the previously visited location and back again respectively.²²

An agent that has accrued a collection of rules can use this knowledge to plan a route to reach a needed resource. To do this, the agent generates a sequence of landmark descriptions using the rules, which will form a path that the agent can follow. The agent starts with the right-hand side of a rule in which the left-hand side is *goto(resource)*, where *resource* is either *food* or *water* depending on the agent's needs. If multiple rules leading to the correct type of resource are known, the rule that leads to the resource which the agent most recently visited is used. From this initial rule, the agent can plan a path by iteratively replacing the *goto(_)* term in the current expression as shown in the example below. In this example, an agent knows the rules given above, needs food and is positioned at *PosY*. It is able to plan a route to food by taking the following steps;

²²Allowing routes to be stored in both directions, even when they have been travelled in only one direction, can be seen as valid in light of experimental findings in rats. Foster and Wilson (2006) have shown that after making a journey, rats replay the path taken through their mind in reverse order, equivalent to the transformation needed to learn the return route.

$goto(food)$	Target resource
$\Rightarrow goto(PosX)$	From rule (24)
$\Rightarrow goto(PosY) L_1 L_2 L_3$	From rule (25)
$\Rightarrow L_1 L_2 L_3$	As $PosY$ is current position

At the first two stages, the $goto(-)$ term is replaced by the right-hand side of a rule which has this term as a left-hand side, until this term becomes $goto(PosZ)$, where $PosZ$ is the last position that the agent visited. The term can then be removed (as the agent is either at this position already, or can get there using the information in its short-term memory) and the remaining expression will be a list consisting of descriptions of landmarks that should be passed in the given order to reach the resource. In the case of several rules being available to replace a $goto(-)$ term at any stage, each of these options will be considered to ensure that the shortest possible route is found.

5.2.3 Language

From the description of the rules used for navigation, it should be apparent that they are actually rules of a regular grammar, in which the $goto(-)$ terms of each rule play the role of non-terminals and landmark descriptions are terminals. Furthermore, the process described above for combining rules to produce a route to a resource is equivalent to generating a sentence of a regular language. Similarly, following a route generated in this way amounts to parsing a sentence of the language. This is important: *if a regular language parser could help navigation, it may first have evolved for this purpose.* A subsequent change in the neural connections might have made this parser available to the human speech circuitry.

Given agents who have the ability to internally store and utilise navigational information in the way outlined above, all that is needed for them to be capable of exchanging routes to a resource location is a shared lexicon of names for landmark properties. Should this exist, an agent can inform another of a route by planning it as if it were to follow the route itself, then articulating this list with each landmark description replaced by

appropriate words.

In this thesis, the lexicon is assumed to be pre-existing, perfectly coordinated and fixed in the population. In supposing that words may pre-date syntactic language, this approach agrees with Bickerton (1998) and Jackendoff (2002), both of which suggest a possible phase in the evolution of the language faculty during which language featured words but no syntax. Steels (1996b) amongst others has shown, through computer simulation, a solution to the lexicon coordination problem. Given that this has shown to be viable, the decision has been taken here not to repeat this work.

When an agent is told about a route by another agent (the situations in which this may occur are outlined in section 5.3), it is stored in the same data structure as those rules created through exploration and as such modifies the listener’s grammar. It is necessary for the listener to add three new rules: the first captures the information that the agent has just been given on how to travel from the point at which the conversation is being held to the resource, while the second and third record the path between the last location the agent visited and the current one, and the path back again. For example, should an agent travel from a location $PosA$ past landmarks L_1 , L_2 and L_3 to another location $PosB$, and here be told that food can be found by following the route ‘ $W_1 W_2 W_3 W_4$ ’,²³ it will create the following rules;

$$goto(food) \Rightarrow goto(PosB) W_1 W_2 W_3 W_4 \quad (27)$$

$$goto(PosB) \Rightarrow goto(PosA) L_1 L_2 L_3 \quad (28)$$

$$goto(PosA) \Rightarrow goto(PosB) L_3 L_2 L_1 \quad (29)$$

While rules 24 and 27 both contain directions to a food resource, a difference exists between them on which it is necessary to comment. In rule 24, the description of the location of the food is present, while in rule 27, a description of the location is only implied from a previous location and a route to be followed. The reason for this dissimilarity in the rule sets is that, in the case of an agent gaining a rule linguistically, it has no knowledge of the location at which the resource is located; as the description of the location is not vocalised by the speaker, the listening agent has no way of acquiring

²³Where W_n is the word used to articulate the landmark described by L_n .

knowledge of the location's appearance until it follows the given route and locates the resource itself.

In addition to this difference, there is one other distinction in the way rules are used in language and navigation. As noted earlier, an agent will always generate a route to the resource of the required kind that it most recently visited, whether the route is required for the agent's own use or is to be communicated. However, if the agent has not visited any resources, it will use rules it has gained through language for its own navigation, but will not use these rules for forming routes to communicate to others: second-hand information is not shared with others. The results of preliminary experiments (Kazakov and Bartlett, 2004a) revealed that were this distinction not made, then information about resource location gained linguistically was frequently unreliable. Due to the exhaustibility of resources, agents were able to potentially pass on routes to resources that were no longer there; in a situation in which agents were able to pass on second-hand knowledge, it was often observed that obsolete information would be recycled between agents, and be difficult to remove from the shared body of knowledge.

This unreliability is unrelated to the issue of trust, an issue sometimes studied in agent populations in which lying is possible. Here, the dynamic nature of resource location leads to agents accidentally passing on out-dated information, rather than doing so through incompetence or with malicious intent.

While the form of language outlined here is undoubtedly basic, it does have several important properties of human language. As the rules form an example of a regular grammar, *recursion* is possible and indeed occurs. Given that the language is recursive, it must also be *infinite*. Finally, the language also exhibits *compositionality*. These issues, and others related to the language type, are explored in detail in section 5.6.

5.3 Agents' Behaviour

A range of agents, utilising a variety of different navigational and linguistic abilities, are needed in order to test the veracity of the claim that syntactic communication may have

evolved from navigational abilities and may have been used to communicate directions to resources. However, the basic behaviour of all agents is very similar and follows the structure laid out in figure 16. The topic of interest for the experiments in this thesis is the advantage that may be brought to a population by the use of syntactic language. Therefore, the actions taken by an agent possessing this abilities will be explained now, with an explanation of how other behaviours differ from this being given in the next chapter. As explained above, the relative merits of these behaviours can be found by comparing the performance of individuals or populations exhibiting them.

The actions taken by all agents in the simulations is based on the value of three *drives*, which describe the level of three internal states: hunger, thirst and sex drive. The former two of these are used to determine the action that each agent will undertake at each timestep, while the latter is used to determine if an agent is willing to mate should the opportunity arise. As the first two drives are used differently to the sex drive, hunger and thirst will first be discussed, then the sex drive considered along with mating behaviour.

At the beginning of each turn, the hunger and thirst drives are increased to represent the cost of living. If the agent's drives exceed the maximum allowed, it is considered to have died and is removed. The sex drive also increases at each timestep, however in contrast to the other drives, agents do not die when the drive reaches its maximum. As some aspects of behaviour such as this, are more simply explained by reference to the amount by which a drive can further increase before it reaches this maximum, the concept of a *resource level*, as a dual to a drive, is therefore also required and defined as follows:

$$\text{resource level} = \text{maximum drive level} - \text{drive level}$$

The action that each agent takes at each timestep is determined by the value of its hunger and thirst drives. More specifically, if a drive level exceeds a particular threshold value, then the agent will carry out an action to attempt to reduce the appropriate drive. If both drives exceed the threshold then attempts will focus on reducing the drive which is greater. The agent attempts a sequence of actions in trying to find a necessary resource, which is shown in figure 16 and is explored in more detail below. If

```
hungry := false;
thirsty := false;
if hunger  $\geq$  thirst and hunger > drive threshold then
    hungry := true;
else if thirst > hunger and thirst > drive threshold then
    thirsty := true;
else if no location of food is known then
    hungry := true;
else if no location of water is known then
    thirsty := true;

if hungry = true then
    if food in current square then
        Eat food;
    else if food in adjacent square then
        Move into square with food;
    else if can generate a route towards food then
        Generate a route and follow it on subsequent turns;
    else if another agent is present then
        Ask for help;
    else
        Make a random move;
else if thirsty = true then
    if water in current square then
        Drink water;
    else if water in adjacent square then
        Move into square with water;
    else if can generate a route towards water then
        Generate a route and follow it on subsequent turns;
    else if another agent is present then
        Ask for help;
    else
        Make a random move;
else
    if at a resource then
        Make a random move;
    Attempt to mate;
```

Figure 16: The behaviour followed by agents in the model.

neither drive exceeds the threshold, no action will be taken to find food or water and the agent will behave differently, including by mating if the opportunity exists.

As noted previously, an agent whose hunger or thirst levels reach their maximum will die of hunger or thirst respectively. Additionally, an agent may die of old age. After an agent has existed for a 300 timesteps, it faces the prospect of dying of old age with constant probability at the beginning of each subsequent timestep.

If an agent is neither hungry or thirsty, it will have the opportunity to mate during its move. In order to do this, it must be sharing a location with another individual, who is also neither hungry or thirsty, and both agents' sex drives must exceed a given threshold. Any such pair of agents may mate: agents do not have an associated gender. Mating costs each parent a third of both of their hunger and thirst resource levels, which are transferred to the child as its initial resource levels. The child is initially placed in the same square as its parents or, should that square be fully occupied by a combination of agents, landmarks and resources, in an adjacent square. All child agents begin with identical behaviour: variation in behaviour between individuals within a population is entirely due to the different knowledge they process (a learned trait), not different ways of processing this knowledge. For this reason, it is unnecessary (and indeed meaningless) to create children through cross-over or mutation. The sex drives of both parent agents are reset to zero after mating.

The remainder of this section now returns to the behaviour to be taken by an agent in attempting to reduce hunger or thirst. As mentioned above, the description given here is specifically that taken by agents who navigate and communicate using the representation given in the previous section. However, the behaviours of other agents can be fairly simply described by how they differ from this behaviour.

If it has been determined that an agent is in need of a resource (i.e. if its drive level exceeds the threshold), the agent attempts to reduce this drive by utilising an appropriate resource. In the case of hunger, this will be a food resource, whereas for thirst it will be water. The agent first looks for a resource in its current location and consumes from it if one exists. If no resource is present, then the agent examines the adjacent

squares, and moves into any of these which contain the desired resource. Consumption will then be possible on the next turn. Consumption from a resource reduces the drive level to zero, and removes one unit from the resource consumed, regardless of the size of the change in the drive level of the agent. When all units of a resource at a particular location are consumed, the resource is removed. Depending on the experimental setting used, a new resource may then appear at another location in the environment or the resource may reappear at its current location after a period of time.

If a resource could not be directly observed, the agent will next attempt to plan a route to a resource. This takes place using the representation and method described in the previous section. Agents retain knowledge of the location of these resources in the form of routes as they explore the environment, and if the agent has previously visited a location containing the appropriate resource, it will now be able to use this knowledge to generate a route which it can then follow to the resource on subsequent turns. If an agent does not know about a resource, it will be able to ask any agent sharing its location if that agent knows of a resource. If it does, then the agent will store the route it has been given by this agent, and then proceed to follow it on subsequent turns.

Finally, if the agent has no other options, it will make a random move in the hope that its new location will place it closer to a resource, or another agent who can direct it to one. In order to ensure that agents extensively explore the environment while searching for resources, their movement is generated randomly, but with a tendency to remain travelling in the same direction. If such a bias is not implemented, the most likely behaviour is to wander back and forth in a small area, continually exploring the same relatively small set of squares. In contrast, the method used has a clear inclination towards entering a square which has not been previously entered. Specifically, an agent has a 50% probability of continuing in the direction in which it was moving, and otherwise deviates right or left by a one square with equal probability. For example, an agent which headed north on its last move, will have a 50% chance of continuing north, 25% of travelling north-east and 25% of proceeding north-west. Examples of this are shown in figure 17. If continuing in the same direction would take the agent into a square which it would not be able to enter, either because it is full or it is outside the

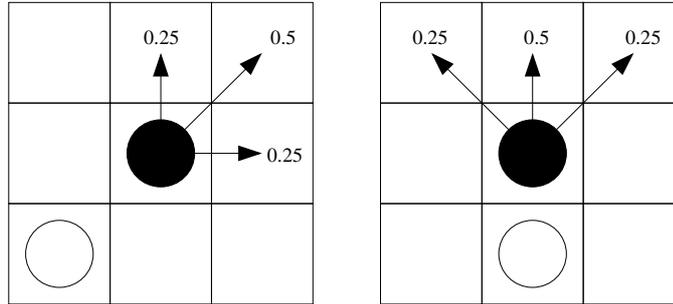


Figure 17: Examples of the forward-biased movement of agents. The agents (shown in black) continue in the direction in which they are heading with probability 0.5, or move one square to the left or right with probability 0.25 each. The position of the agent at the previous step is shown as a white circle.

edge of an environment, it will instead turn left or right with equal probability until a square is found which it can enter.

Clearly the outlined behaviour for attempting to locate resources leads agents to become reliant on a small number of resources at any one time rather than fully exploring the environment to locate multiple locations at which they may feed. Exploration is a secondary task, which is only undertaken when entirely necessary. In this aspect, the agents' behaviour may be seen as analogous to that of several species in the real world, for example ants, which search for resources which they then fully deplete before searching for the next.

5.4 Experimental Details

From the higher-level, conceptual explanation of the model given above, this chapter now proceeds to a more detailed description of the constructed simulation including many details which are irrelevant to the understanding of the concepts that this model addresses, but may prove to have some bearing on the outcome, and are certainly required if the experiments performed are to be reproducible. The majority of the implementation details of the simulation are undoubtedly unimportant and much the same outcomes would probably be obtained should different design decisions have been made, however this follows the approach adopted by Ackley and Littman (1994);

We try to be as complete about details of the model as we can, not because we believe every detail to be critical — some of them certainly seem minor — but to be clear “up front”, and because, as experience with computer models of complex systems shows over and over again, it is nearly impossible to be certain *which* details don’t matter.

5.4.1 The Environment

The environments consist of a two-dimensional grid of squares arranged orthogonally. Each of these squares has room for 4 entities, which may be agents, resources or landmarks. The layout of the grid is such that each square (except those at an edge) is considered to have 8 neighbouring squares, including those diagonally adjacent, as well as those vertically and horizontally adjacent. Agents can see the contents of each of the neighbouring squares, and when moving, may move to any which contains less than four entities. There is no ‘wrap around’ at any edge of the environment: the perimeter of the environment functions a wall through which agents may not pass. Environments of several different sizes are utilised, this being one of the parameters studied in section 6. The environments studied range between 5 and 13 squares in length and width.

Other than agents, environments feature two immobile entity types; resources and landmarks. Each resource has an associated type, food or water, and a resource level, corresponding to the number of ‘meals’ it has left to provide before becoming exhausted. When a resource becomes exhausted, it is removed from the environment. The strategy subsequently used to add additional resources to an environment after some have become exhausted is a parameter of the experiment undertaken, as is the initial resource level. A resource occupies a single unit of capacity in an environment square, regardless of the number of ‘meals’ that it still contains.

Landmarks are stationary items which are regularly spaced throughout the environment to provide markers by which to navigate. In order for the routes that agents gather and exchange to be meaningful, it is necessary for adjacent landmarks to be visible from the square containing a landmark. To this end, agents have the ability to observe

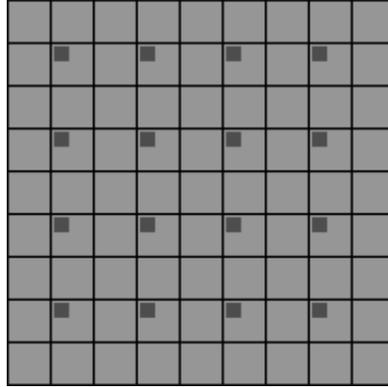


Figure 18: The configuration of landmarks used in experiments. Landmarks are shown as boxes distributed regularly through the environment.

landmarks up to two squares away, including diagonally, and landmarks are arranged in an orthogonal grid with adjacent landmarks separated by two squares, as shown in figure 18. This positioning allows an agent to always observe between four and nine landmarks depending on their position. When determining a location for use in navigation, it is this set of observable landmarks that an agent records. As well as being necessary for landmarks to be readily visible, it is necessary for them to be distinguishable from each other. Landmarks therefore have three observable properties (height, weight and colour), each of which has one of three possible values. To further ensure distinctiveness, the properties of landmarks are chosen so that from all positions in the environment, the landmarks observable can all be discriminated between.

5.4.2 Navigation and Communication

Apart from the set of grammar rules that an agent processes, it also retains a short-term memory consisting of the last location it visited, and an ordered list of landmarks that it has passed since then. This knowledge is used in constructing new grammar rules when an agent reaches a location containing a resource, or at which it receives a route from another agent. The location that is held in the short-term memory is actually the location at which the last new grammar rule was created: the short-term memory retains a record of the agent's movement since it last had cause to remember a

route. The functioning of the short-term memory is in fact more complex than simply accumulating a list of landmarks. When the agent moves from the vicinity of one landmark to another, rather than just storing the properties of the new landmark, the short-term memory checks whether this landmark was the one seen previous to the last one observed. If this is the case, the agent must have gone in a loop from the environs of one landmark to that of another and straight back again. In this case, the short-term memory is amended to remove the record of the detour. This has the effect of shortening the eventual grammar rules and thus the generated routes without affecting the agent's ability to retrace the path. Apparent loops with a greater length are not detected or removed as the seemingly identical descriptions may refer to different landmarks with the same perceptual properties. A similar mechanism is used, for the same reason, to eliminate analogous loops from any generated routes.

Like the short-term memory, the path-planning mechanism has complexities which were not mentioned in the earlier section. Amongst these is the issue of how the shortest path to a given resource can be found. This is achieved through an expansion process which generates a tree of (partial) routes which is incrementally expanded using an A* search (Hart et al., 1968) to explore all possible expansions. The metric for this search combines the length of the partial path already found from the destination, which is measured through the number of landmark descriptions in the expansion, and the shortest possible distance from the intermediate position mentioned in the *goto(-)* term in the expression to the location stored in the agent's short-term memory. This guarantees the shortest path will be the first found.

Another factor of route planning neglected earlier is the use of the short-term memory in conjunction with a generated route to form a meaningful path. Informally, the method used is to plan a route which equates with back-tracking along the route held in the short-term memory and then following a route from this location to the target resource. More technically, the process used is to form a route to the required resource by generating a sentence consisting solely of landmark descriptions, from the union of

the grammar that the agent possesses and the rule

$$\textit{goto}(\textit{Pos}X) \Rightarrow L_1 L_2 L_3 \dots L_n \quad (30)$$

where $\textit{Pos}X$ is the location stored in the agent's short-term memory and $L_1 \dots L_n$ is the list of landmarks that the short-term memory contains in reverse order. As rule 30 will be the only one in the grammar without a non-terminal on its right hand side, this rule is invariably used as part of the resulting generation.

One final aspect of the grammar requiring additional comment is the way in which information is removed from the agents' knowledge. While discovering or being informed of a resource location requires an agent to add a rule to its internal grammar, finding that a resource has been exhausted, and thus is no longer available, requires the removal of a rule. Should an agent find itself in the position of having followed a route it has generated, but no resource being available on arrival, it will forget the rule that mentioned that a resource could be found at that location. It will not however remove the rules stating how to get to and from the position in question. These rules may still be needed in forming a path that passes through this location on the way to another resource. As importantly, the agent will probably be in urgent need of these rules in trying to plan a route from this location to another resource to service its pressing need.

5.5 Model Viability

Up to this point, this chapter has described the model that will be used to evaluate the benefits that syntactic language brings. While the next chapter will explore the range of environmental conditions in which syntactic communication proves beneficial, it is first necessary to demonstrate that this model is capable of producing a stable population which implements the behaviour described above. That is to say, before the performance of differing populations can be compared, it must first be established that the model described is capable of producing at least one environment in which agents can survive and reproduce.

To do so, a single experiment is presented here in which a population of agents, who

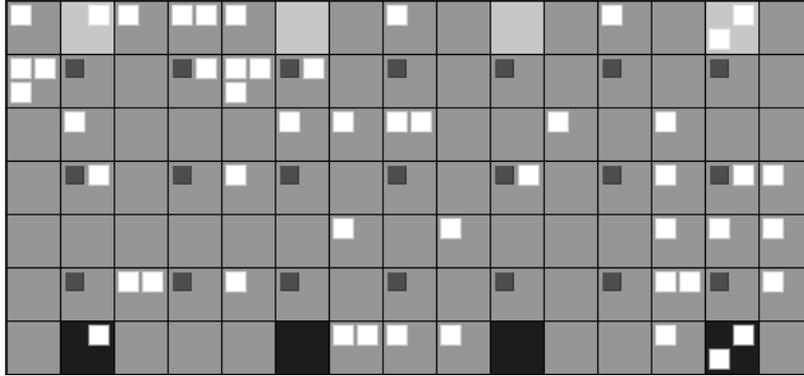


Figure 19: Environment for establishing the validity of the model. The environment contains several food resources, shown as light-coloured squares at the top of the environment, and water resources, which are shown as dark squares at the bottom. Landmarks appear as regularly spaced dark boxes within the environment. Agents are shown as white boxes.

implement the behaviour given above, are placed into an environment. This is then run for a period of equivalent to several generations over which time the size of the population is recorded, as per the definition of the experimental setting given. This is repeated thirty times in order to establish a higher degree of certainty in the result. The environment in which this experiment is carried out is shown in figure 19. This environment is based on a layout style which will be more thoroughly explored in the subsequent chapter. At the present time, it is sufficient to note that the environment contains a regular grid of landmarks for navigation, and several food and water resources, which are spaced out evenly across the top and bottom of the environment respectively. When resources are exhausted, they disappear for ten turns before reappearing at the same location. Each run of the environment begins with 50 agents, who are randomly placed within the environment.

The results from this experimental run are shown in figure 20. From its initial size of 50, the population first increases in size rapidly as those agents who have been fortunate to find both resources quickly begin breeding. This is then followed by a significant drop in population size after 60 cycles, corresponding to the time at which those agents who neither found the necessary resources nor received directions to them begin to die of hunger or thirst. The population then begins increasing in size again before experiencing another decline. This decline begins around the 300th timestep

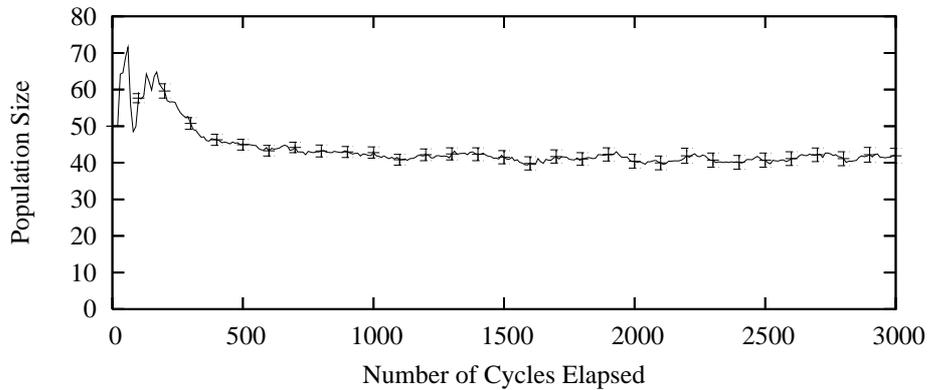


Figure 20: Results of the experiment using the environment shown in figure 19.

when the surviving agents from the initial population (and those born within the first few timesteps) begin to die of old age. As these large decreases in population size at 60 and 300 cycles are due to the fact that a large number of agents are all created at the same time with identical resource levels, no further such effects are seen in the population size for the remainder of the simulation run: after the first generation, the population consists of a group of agents with a wide range of ages and resource levels. Finally after the second of these declines, the population reaches a level of between 40 and 45 before the thousandth cycle, and remains stable at this size for the rest of the simulation.

These results demonstrate that the model allows for stable populations to emerge, which navigate by routes of landmarks to return to locations they have previously visited, and use these same routes to communicate directions between themselves. Further experiments using the same methodology are presented in chapter 6. In these experiments, the aim will be the establishment of the role that the environment has on the population's performance. Furthermore, comparison will be made between the performance of agents implementing the behaviour explained in this chapter with other possible behaviours.

5.6 Theoretic Aspects of the Language

The remainder of this chapter now returns to the representation used for navigation and communication which is at the heart of this thesis. Specifically, the intention here is to illustrate that what appears to be a simple form of communication in fact possesses several key features of human language.

As was explained earlier, the representation is equivalent to a regular grammar. This makes it possible to incorporate recursion, which Hauser et al. (2002) believe to be one of the unique features of human language. The specific form of recursion available in this case is called *left recursion*. In such a grammar, it is possible, from some non-terminal, to derive an intermediate stage in the parsing at which the same non-terminal appears as the leftmost symbol of the intermediate form.²⁴ In the case of the grammar used here, this means that after a number of steps, a path to *PosZ* is described in terms of going to *PosZ* and then following a route. Recursion within the grammar occurs in the most simple case in those rules directing the agent to a point and back again, for example in rules 25 and 26 given earlier;

$$\begin{array}{ll}
 \text{goto}(\text{PosX}) & \text{Right hand side of rule (25)} \\
 \Rightarrow \text{goto}(\text{PosY}) L_1 L_2 L_3 & \text{Left hand side of rule (25)} \\
 \Rightarrow \text{goto}(\text{PosX}) L_3 L_2 L_1 L_1 L_2 L_3 & \text{Using rule (26)}
 \end{array}$$

More generally, recursion occurs in any set of rules which define a path incorporating a complete loop. In the grammars that emerge from the model, most recursion will be *indirect* recursion as shown above, that is to say, recursion in which at least two rewrite steps are needed to form a recursive definition. It is possible however, for immediate recursion to occur, if an agent discovers a circuit from a location past a series of landmarks and returning to the initial position.

²⁴It is relatively simple to convert the representation used in this thesis into a *right recursive* grammar. To do this, rules are rewritten from the current form, which are equivalent to stating ‘to go to X, go to Y then follow the given path’, to a form equivalent to ‘starting at Y then following the given path will result in going to X’. The choice of which one of these forms is used has no bearing on the complexity of the language learnt, but has implications for whether top-down or bottom-up parsing can be used.

As the representation allows recursion, it will also allow for an infinite number of strings to be produced. The infinite size of human languages are taken to be another important distinction from animal communication systems, which are typically fixed in size to a small number. It can be easily seen that the language produced must be infinite as any agent knowing the route to a location and back can produce an infinite number of paths to that location using only those two rules. The first path goes directly to the location; the second one goes to the location, back to the current position and then back to the target; the third goes to the location, back to the current position, back to the target location, back to the current position and then finally to the target again. An infinite number of routes can thus be generated by adding an additional loop to the end of the route each time. Any agent possessing more than the two rules used here can clearly still produce an infinite number of paths.

One final property possessed by the language is compositionality. To be truly compositional, and not merely syntactic, the meaning of a sentence must be a product of the meanings of the parts of the sentence and the way that they are combined. In the language type described here, this is clearly the case; a sentence as a whole has a meaning, as does each part of it (the landmark to which it refers), and a sentence's meaning depends on the order in which each of these parts is given, not just which ones are present. Furthermore, if landmarks are permitted to have non-unique descriptions, then the landmark referred to by a given description may be different depending on the tokens before and after it, in a similar way to how the parts of a sentence in human language can have their meaning modified by the context in which they appear.

The representation used in this thesis is equivalent to a regular language, however modern human languages are generally taken to be (at least) context-free. As the stated aim of the thesis is to demonstrate that human language may have evolved from the earlier navigational abilities already present in humans, consideration must be given of how a context-free language could have evolved based on a navigational task requiring only regular grammar manipulation. Two explanations for this are possible. Firstly, a regular language parser could have been exapted from its navigational role and subsequently have evolved context-free abilities under some pressure arising from

language use. An exploration of the linguistic requirements that may have made this possible are outside the scope of this thesis, which is concerned with the link between navigation and language, not the pressures which may have acted on language use.

A second explanation, which is within the scope of this thesis, is that abilities equivalent to using a context-free grammar may have been needed for some navigational task. A task that requires context-free complexity has been outlined in a previous paper (Kazakov and Bartlett, 2005), and is that of going to a location and then returning along the same path. This is a seemingly simple activity which one would expect most creatures to be capable of, yet it is impossible to specify using only a regular language representation for navigation. If one returns along the same path as one takes, the list of landmarks passed along the path travelled will be palindromic; it is a well-known result that a context-free grammar is needed to recognise palindromes. If the ability to represent and handle context-free grammars were needed to return along a path while navigating, it may have been possible for this ability to be incorporated into the speech mechanisms, by what may have been a relatively small mutation, to allow context-free language.

5.7 Summary

This chapter has described a model which has been developed to test the plausibility of the claim that syntax in language may have evolved from an analogous ability used for navigation. At the heart of the presented model is a system of knowledge representation, which is based on the routes used to navigate by animals. The model makes a dual use of this representation for use in both the navigational task and in communicating directions to others. The regular language that this permits, while basic, possesses several key characteristics of human language, specifically recursion, infinite size and compositionality. A communication system similar the one used appears relatively easy to evolve by incorporating the syntax handling abilities that are inherent in the navigational system into a more primitive system of lexical communication. A discussion of how this language may be extended to become context-free was given.

Through the use of an experimental run of the model, it was shown that the navigational and communicational methods used are capable of generating a population with a stable size. This shows that the model is a viable method of behaviour for agents. In the next chapter, this experimentation is extended to establish under which environmental circumstances the behaviour outlined in this chapter is most useful. Following this, experimentation is presented in chapter 7 which investigates the evolutionary pathway by which the communicative behaviour may have arisen from its navigational analogue.

6 The Role of the Environment

The previous chapter of this thesis outlined a model for the study of communication using a simulated environment. The discussion of the internal representation used in this model, in conjunction with earlier chapters, discharged the first objective of this thesis, that of demonstrating a link between the abilities needed to process navigational instructions and those needed for syntax within communication. This chapter now continues the task begun in the previous chapter of tackling the thesis's second objective, explaining how such a syntactic language capability could spread through a population by the evolutionary benefits it provides.

Taking from the previous chapter both the supplemental theory that early language was used to communicate directions to essential resources and the model which was developed from this, this chapter and the subsequent one present experimental results illustrating the benefits of language use to a population. The major challenge to agents in the model is that of locating resources which they can then use to survive. As the developed model is based on the theory that an early use of language was to share directions to these resources, it is clear that the environment in which this communication occurs is likely to have a large impact on the utility of this communication. Intuitively, the easier it is to find resources, the less useful the ability to acquire information from others becomes. Likewise, the more 'sensorimotor toil' (Cangelosi and Harnad, 2000) is required to locate these resources, the more appealing it becomes to procure the knowledge linguistically.

The investigation into the impact of environment on the resulting population forms two parts. In the first, the effect of the environment on the resulting size of a population implementing the behaviour outlined in the previous chapter is shown. Following this, the second part utilises control populations to establish not only the influence of the environment relative to other environments, but also relative to populations unable to communicate. The results from the second part demonstrate that, as well as having an influence on the size of the population that results, the environment also plays a

significant role in determining the relative merits of different behaviour. In both parts of this experiment, two aspects of environmental design are considered; resource volatility and resource positioning. In each case, the relative merits of each strategy are shown to be dependant on the environmental factor studied.²⁵

This rest of this chapter proceeds as follows. The aspects of the environment which will be assessed for their impact on the populations are introduced and explained in section 6.1. Following this, section 6.2 shows the effect that these parameters have on the performance of the population. Finally, the *relative benefit* of language use within the context of these environments is found through comparison of the performance of a language using population with controls implementing similar, but non-linguistic, behaviours.

6.1 Environmental Factors Assessed

A great many aspects of the environment could be assessed for their possible impact on population size. This thesis, however, restricts itself to the examination of two properties, both related to the resources within the environment. However, the methodology used to study these issues readily generalises to allow the study of many other environment factors which may be of interest. The first of the properties studied is the positioning of resources, or to be more precise, the spacing between resources. The second property is that of the volatility of resources. In both settings, as many other factors as possible are held constant while the ones of interest are varied.²⁶

²⁵Some of the results presented in this chapter have been previously published elsewhere. Those experiments concerned with resource spacing appeared in Bartlett and Kazakov (2005a), while those studying resource volatility appeared in Bartlett and Kazakov (2005b).

²⁶It is not always possible for all factors not being studied to be kept constant. For example, in the resource spacing environments which will be presented, it is impossible to vary the spacing between resources without simultaneously altering either the size of the environment, the number of resources, the proportion of squares containing resources or the geometric relationship between resource locations.

6.1.1 Resource Spacing

Perhaps the most obvious factor which one would expect to affect the performance of a population is the availability of resources, that is to say, the degree to which resources are likely to be found through exploration. The greater ease with which resources can be found, the better a population would be expected to perform. However, one would also expect the ability to communicate directions to be less important to population success in situations in which exploration led reliably to the discovery of resources.

Several schemes exist which could be used to assess the effect of resource availability, the most obvious of which is to vary the number of resources in the environment. However, the approach adopted here is to maintain the same number of resources but vary the distance between resources. Controlling availability in this way has several desirable properties which will be explored below. As there are two types of resource (food and water), there are two spacing parameters that can be manipulated, the distance between resources of the same type, d_{same} , and the distance between resources of different types, d_{diff} .²⁷ In the following experiments, both of these parameters are varied over a range of values to assess the influence of both independently.

The basic design used for experiments involving resource availability is shown in figure 21. In this environment, food resources are placed in a non-contiguous row at the top of the environment, with water resources placed likewise at the bottom. To alter the values of d_{same} and d_{diff} , the environment size is changed. In each case, the width of the environment is chosen so that four resources can be placed along the length of the environment with a gap of d_{same} between them, and one square space between the end resources and the environment sides. Likewise, environment height is altered to allow d_{diff} squares to fit between the resources of different types. A range of environments formed by altering the values of d_{same} and d_{diff} are shown in figure 22.

²⁷In fact, it is possible to use three parameters, the distance between water resources, the distance between food resources and the distance between resources of different types. As this necessitates more complex environment design than that which is presented here, the distance between food resources and the distance between water resources are kept equal in this work.

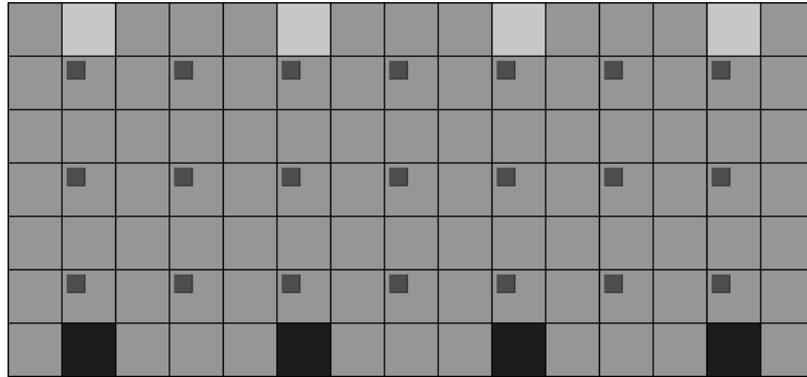


Figure 21: Typical environment layout for the resource spacing experiment. Environment shown has $d_{same} = 3$ and $d_{diff} = 5$. For key, see figure 19.

In choosing to assess the effect of resource availability by altering spacing between resources rather than the more intuitive method of varying the number of resources in the environment, it proves impossible to hold another parameter constant, that of environment size. This scheme is still preferable to that of the alternative however. Altering the number of resources in the environment has two undesirable properties. Firstly, increasing the number of resources would reduce the distance between resources. As the experiments presented later show resource spacing to be a very important parameter, this is very disagreeable. Workarounds to this problem can be found by altering the environment size or through the arrangement of landmarks, however a second problem is insurmountable. As a maximum of three agents can be in a cell with a resource at once, introducing more resources increases the number of agents who can feed (or drink) at each timestep. As the carrying capacity of the environment depends on the rate at which agents can acquire resources, altering the amount that can be theoretically taken each turn clearly has very serious implications. In contrast, the chosen scheme retains four locations from which each type of resource can be acquired for all parameter settings.

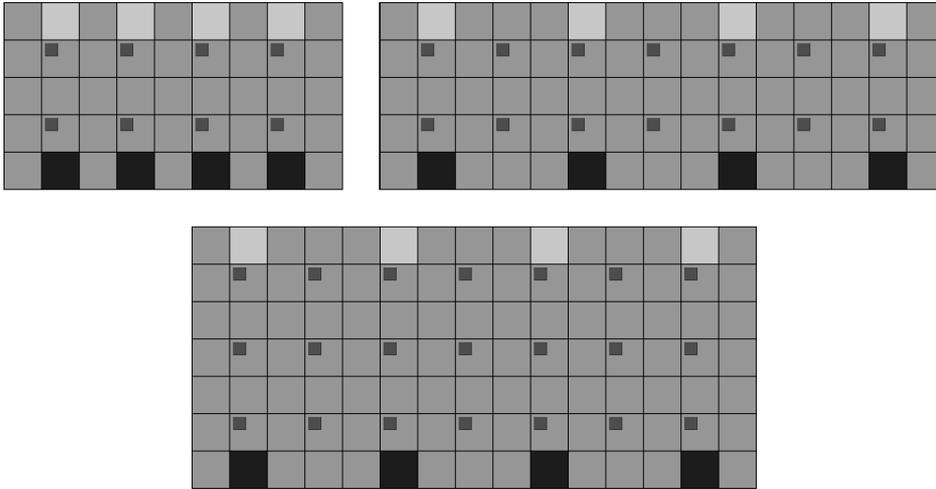


Figure 22: Various environments used for resource spacing experiments. Each environment has different combinations of d_{same} and d_{diff} . Top left: $d_{same} = 1$, $d_{diff} = 3$. Top right: $d_{same} = 3$, $d_{diff} = 3$. Bottom: $d_{same} = 3$, $d_{diff} = 5$. For key, see figure 19.

6.1.2 Resource Volatility

The other environmental factor considered in this thesis is the volatility of the resources. One reason for considering this factor from amongst all those possible is that it is completely independent of the resource spacing factors: while studying the spacing of resources, volatility need not be considered and, likewise, a single set of spacing parameters can be used while the factors determining volatility are varied. Most other interesting factors which could be studied involve moving resources in the environment, overlapping with the spacing studies.

If the purpose of language was for sharing directions, the volatility of the resources to which directions were given would intuitively be expected to have an impact on any benefits related to language use. Communicating directions to entirely static resources should prove more advantageous to a listener than if the resource is highly volatile and may no longer be found at the communicated location. However in contrast, the speaker would lose less in telling another about a resource that was likely to have been exhausted before it could return to it again than in informing another individual about a resource that would almost certainly remain in place to satisfy its own future needs.

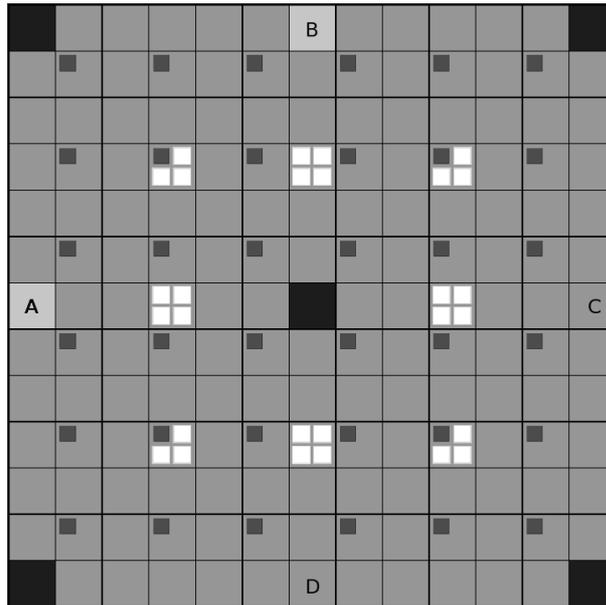


Figure 23: The environment layout used for volatility experiments. When food resources are exhausted, they are replaced by new resources at the geometrically opposite point in the environment. (i.e. resources move $A \Leftrightarrow C$ and $B \Leftrightarrow D$). Water resources are static and inexhaustible. For key, see figure 19.

In contrast to the resource spacing experiment, all experiments assessing the impact of resource volatility utilise an identical environment map. The initial configuration of this environment is shown in figure 23. The water resources are made static and inexhaustible, while food resources are volatile, and when exhausted are removed from the environment and replaced immediately with an identical resource at the opposite side of the environment.²⁸

As with spacing, two factors contribute to the aspect under consideration; *consumption* and *decay*. In the previously presented model, resources become depleted by use and are exhausted when a given number of ‘meals’ have been consumed from them. Altering the number of meals that can be provided before exhaustion affects volatility: the more meals that a resource is initialised with, the less volatile it will be. The

²⁸The new resources are always placed geometrically opposite the exhausted ones, rather than randomly, in order that distances between resources remains constant. Varying distances randomly at the same time as studying resource volatility would have an unpredictable effect to the population dynamics as both these factors affect the relative merits of the differing strategies.

second factor, decay, is introduced by altering the model so that the level of each food resource decreases by a constant amount each turn, in addition to any decrease due to consumption. The larger this decrease, the more volatile the resource. Both these factors are studied over an exponential scale over a large interval to allow effects at both large values and small values to be observed. Specifically, both parameters are varied between 25 and 3200, doubling in value at each step.

At this point it should be noted that the amount of food available in each environment is constant and in no way dependant on either factor of volatility. If this were not the case, while it would be valid to compare performances of different populations in the same environment, comparisons between behaviour observed in different environments would be meaningless. However, it is actually the case that the immediate replacement of one food resource by another when depleted ensures that two food resources are present at all times, allowing six agents to eat during each turn. The only factor varying across experiments is the rate at which these two resources change position.

6.2 The Effect of the Environment on Benefits

In seeking to understand the role that the environment may play in determining the usefulness of communication, it must first be established that the environmental factors chosen for study impact upon the resulting size of the population. In doing so, the aim is to determine only how the population size depends on these factors, not how beneficial the implemented behaviour is relative to other strategies.

For reasons which will become clear later, the quantity that is used as a measure of the population's success is the average population size over the course of the simulation. The more intuitive measure of the final population size is rejected because not all populations tend towards a stable size, but may fluctuate in magnitude. This is seen below in the experiment involving resource volatility. If the final population size is used in such cases, the point at which the simulation ends may significantly alter the perceived success in the given environment. In contrast, utilising an average size

avoids this problem. For populations which have a stable long-term size, the average population size will tend to this size as the length of simulation is increased. Similarly, for populations not reaching a stable size, the average population size tends to the value about which the population size fluctuates.

6.2.1 Resource Spacing

To obtain data on the effect of the spacing of resources on the success of the population, the procedure used in section 5.5 is used. However, whereas in that section the population was tested in only a single environment, here a range of environments are used. These environments have differently spaced resources as explained in section 6.1. Both of the parameters affecting the resource spacing, d_{same} and d_{diff} , are varied over a range of values in order that their effect can be seen.²⁹ For each environment, a series of runs are performed with agents beginning at random positions each time.

By recording the population size throughout an experimental run, it is possible to obtain two forms of graphical output which together illustrate the effect of a given parameter. Firstly, the average population size over the course of the experiment can be calculated, and illustrated as a three dimensional surface plot. In such a graph, the average size is plotted against the two parameters being studied, d_{same} and d_{diff} in this case, and decay and resource size in the case of the volatility experiments. Secondly, the data can be plotted as a time series graph, showing how the population size changes over time. This can allow important information about the temporal dynamics of the population to be seen.

The 3-dimensional graph obtained for the resource spacing experiments is shown in figure 24. This graph shows both d_{same} and d_{diff} to have an effect on the resulting population size, though with d_{diff} substantially more important. As the space between

²⁹Only odd values of d_{same} and d_{diff} are considered here. Due to the arrangement of landmarks, populations perform differently when placed in environments with odd or even values of these distances. The effects of varying these parameters are the same for even values as those reported for odd values, though the size of the populations that result are different.

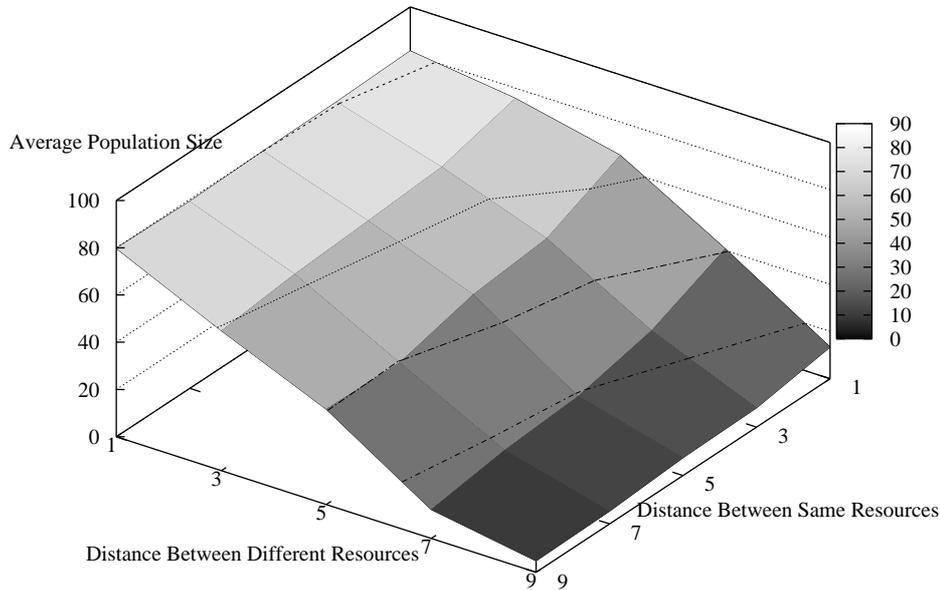


Figure 24: The role of resource spacing in determining population size. All results are averaged over 30 runs and 3000 time steps.

food and water is increased, the average population size falls dramatically. Altering the space between resources of the same kind has a similar effect, but of a smaller order of magnitude. These results make intuitive sense; the closer resources are together, the easier it is to find them, and therefore the less agents dying of hunger or thirst. The larger role played by d_{diff} can also be explained. As an agent must repeatedly gather food and water, it will make a substantial number of journeys between food and water resources. This makes d_{diff} an important parameter. It is only on the occasions that a food or water resource is unavailable due to exhaustion that d_{same} becomes important. Given the relative infrequency of this, d_{same} has a smaller role to play.

While it is possible to plot the changes in population size over time for each point studied, such analysis would be very space consuming and reveal little of interest. Instead, such results are presented for a few points to illustrate the broader situation. Figure 25 shows the population's behaviour at four points in the studied parameter space. The graphs shown, from top to bottom, move progressively from the back corner to the front corner of the graph in figure 24.

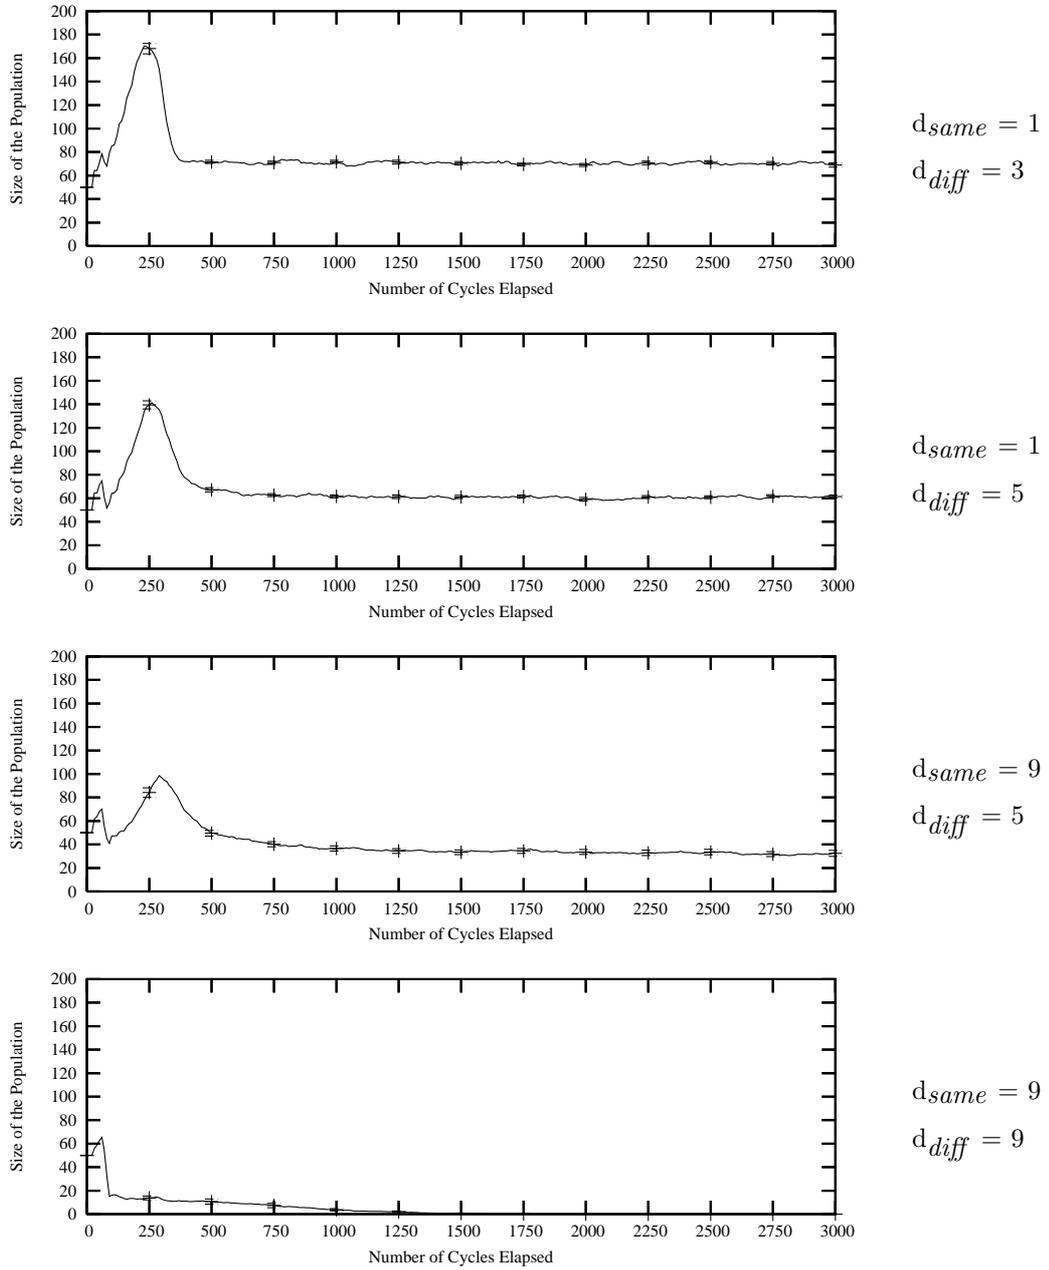


Figure 25: Comparison of population dynamics for different values of d_{same} and d_{diff} . All results are averaged over 30 runs. Note the high degree of consistency between runs in this figure, and all others, as indicated by the extremely small standard errors shown.

Through these graphs, a common trend emerges. First the population increase in size for a few cycles, before decreasing again at the point at which agents who have failed to find resources die. Following this, a second phase of growth begins, after which the population decreases to reach the level it maintains for the rest of the simulation. If the initial drop in population size is great enough, no second wave of increase/decrease occurs and the population heads towards extinction. The second decrease occurs from around 300 cycles, and is due to the death from old age of all those agents who remain from the initial population. As populations past this point consist of agents with a broad mixture of ages, such a synchronised decrease is not seen again.

The smaller d_{same} and d_{diff} , the larger the increase in the population size is seen. Similarly, the smaller these values, the smaller the drop in population size at the two large decreases. The ramification of these dependencies is that the smaller d_{same} and d_{diff} are, the larger the final population size.

6.2.2 Resource Volatility

Following a similar method to that used for the resource spacing experiment, the effect of resource volatility on the benefits associated with communication can be measured. The environment shown in figure 23 is utilised as described in section 6.1. Whereas agents were randomly distributed in the resource spacing experiment, the fact that the environment size and shape is unaltered between experiments in this environment allows for the agents to begin in the same place in each run. Agents are shown at these starting locations in the figure.

In studying the impact of volatile resources, it proves necessary to alter the threshold at which agents became hungry and thirsty. While it would be desirable to use the same parameters for both experiments, the different difficulties experienced in the two experiments require different thresholds if the results obtained are to show meaningful results. In the volatile setting, following routes to previously visited resources frequently results in finding the resource to have been exhausted, and thus a new resource must be found. As this requires a not inconsiderable period of search, the agents must be given

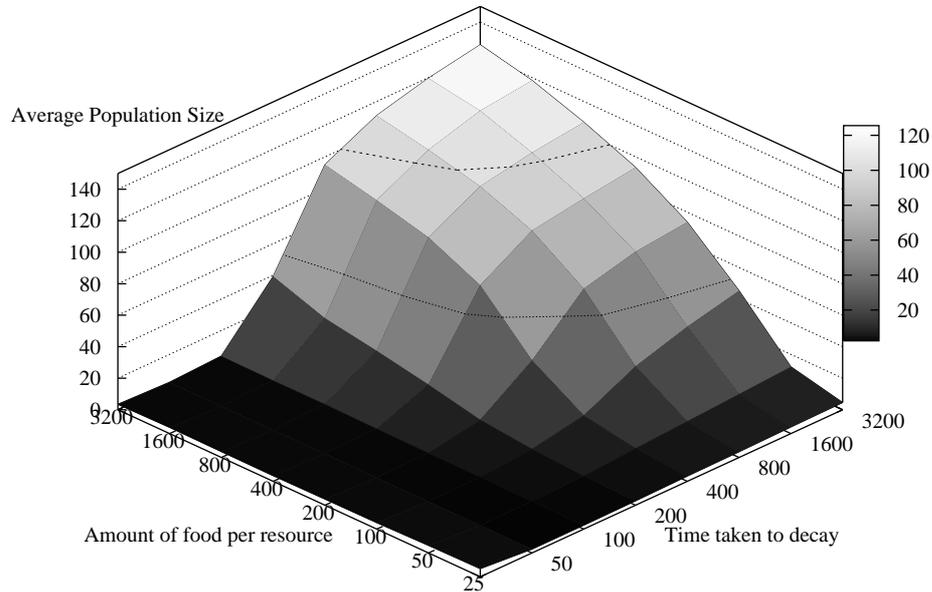


Figure 26: The role of volatility in determining population size. All results are averaged over 30 runs and 3000 time steps.

a relatively long period between becoming hungry and dying of hunger: the threshold value must be set to enable a large number of cycles in which to find food. In contrast, if the population in the resource spacing study is given a long period between needing a resource and dying, the task of finding the resources in time becomes much easier. As this must be done relatively infrequently in that setting, the impact of d_{same} and d_{diff} is significantly decreased. For this reason, lower threshold values are used in studying volatility than were used for studying spacing.

The results of the volatility study are shown in figure 26. The surface obtained fits with the intuitive expectation of the effect of volatility on population size: the less volatile the resources, the larger the average population size becomes. For highly volatile resources, remembering a route to a resource is of little value, as the resource is likely to have moved before it can be returned to. Communicating the location to another agent is unlikely to allow that agent to benefit from the knowledge for the same reason. As resources become less volatile, the ability to return to a previously discovered location becomes increasingly more useful, as does learning a route to such a location from

another.

As with the resource spacing experiment, the population size can be plotted over time to enable the temporal aspects of the parameters' effects to be seen. Figure 27 shows the population size over time, plotted for various points in the parameter space. The graph showing the situation with the most volatile resources adds little to the data shown in figure 26. As the data in this first figure implies, the population quickly declines in size and dies out. However, the other graphs presented show more interesting behaviour.

The most interesting feature seen in these graphs is the periodic increase and decrease in population size.³⁰ This is due to a fall in the population size when a resource moves to a new location and many agents starve before they manage to find the new resource. After the resource has been located again by sufficient agents, the population size begins to increase as fit agents are again able to meet and breed. When the resources are very static, as in the bottom graph of figure 27, the population manages to reach a stable population size between these increases and decreases. However, when resources are more volatile, as in the second graph, the population never reaches this stable level and instead exhibits a wave-like behaviour.

In the third graph of figure 27, the periodic behaviour witnessed in other graphs appears to be missing. In fact, examination of the data from the individual runs of the simulation reveals that periodic increases and decreases in population size do occur. The reason for these not being visible on the graph is due to the parameters of volatility at that point. As the decay period is long, but the amount of food is low, the overwhelming reason for resources becoming exhausted is due to their use by agents. As the rate at which this occurs will vary slightly between runs, the wave forms become out of phase with each other and cancel out, with the result that the cyclic behaviour is not visible in the average value. In contrast, in the second graph (for which the average population size is approximately equal to that in the third), decay is fast and each resource contains a large amount of food. This leads to decay being the prime determiner of the time at which resources move. As this is constant, the movements remain locked in phase

³⁰The reason for the absence of this behaviour in the third graph will be addressed later.

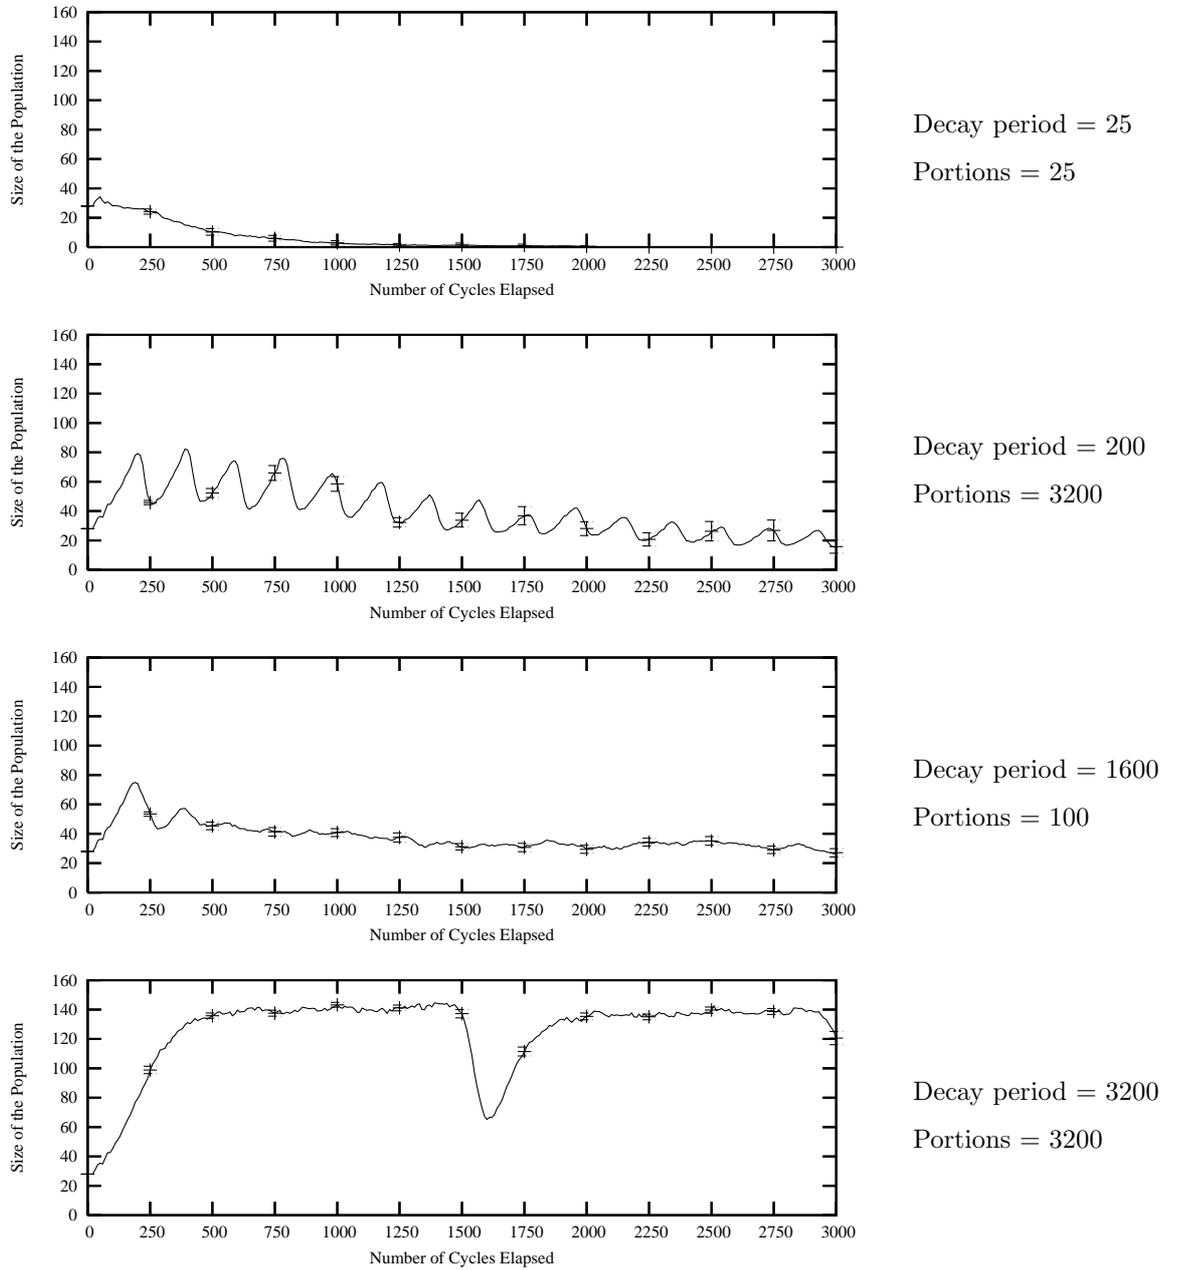


Figure 27: Comparison of population dynamics for different parameters of volatility. All results are averaged over 30 runs.

between runs, and hence are visible in the averaged value plotted.

6.3 The Effect of the Environment on Relative Benefits

In the previous section, the role of the several aspects of the environment in determining the success of populations implementing the communicative behaviour developed in section 5 was examined. While such results may be interesting, they do not in themselves shed light on the major topic of this chapter, that of establishing how such behaviour may have been promoted by natural selection. This goal can be achieved by comparing the data of the performance of populations obtained in the previous section with similar data for other agents who act identically with the exception of not communicating. If, in a particular environment, communication is shown to be beneficial, one would expect that natural selection would act so as to spread this ability through a population.

In order to assess if one behaviour outperforms another, a measure is needed that compares the results obtained from a simulation for a population implementing one behaviour with results for a population implementing another. The concept of relative benefit can be defined as follows:

$$\text{relative benefit}_{AB} = \frac{\text{average population size}_A}{\text{average population size}_B}$$

If $\text{relative benefit}_{AB}$ is equal to 1, then the behaviour implemented by populations A and B can be taken to be of equal benefit to a population. A $\text{relative benefit}_{AB}$ value greater than 1 indicates that the behaviour taken by population A is better than that taken by B , with the greater the value, the more advantageous the behaviour is. Likewise, a value less than 1 indicates B to have the superior behaviour, with smaller values indicating greater benefit.

An alternative formulation would be to compare the final population sizes instead of the average. This, however, is problematic in the volatile resource setting. The results

of this experiment presented in figure 27 show population size to periodically increase and decrease over time. Comparing population sizes at a single point in time rather than taking an average in this setting would potentially lead to a wildly fluctuating value. In contrast, while the definition used does change to some extent depending on the end point chosen, it is far more quiescent.

In addition to assessing the benefit of agents relative to a population who do not communicate, it is also necessary to introduce a second control population to attempt to correct for another aspect of the agent's behaviour, altruism.

The act of sharing routes to a resource is inherently altruistic: the listener gains knowledge of the location of the resource, while the speaker gives up information obtained through sensorimotor toil. Furthermore, the degree of altruism varies when dealing with volatile resources as outlined above, such that the more volatile a resource is, the less altruistic informing another agent of its location is. For these reasons, it is also necessary to obtain results for the performance of a population who are unable to communicate but act altruistically by sharing their reserves of resources instead of directions to them. Comparison of the relative success of the population using language altruistically and those engaged in non-linguistic altruism allows assessment of the extent to which any benefit associated with language use is due to altruism and the extent to which it is due solely to language use.

The non-communicating altruistic agents used are programmed to act similarly to the other non-communicating agents, however if the agent requires a resource and can neither see it or plan a route to it, it may ask any agent sharing its location to altruistically share their stores of this resource. This help is always given if possible and consists of transferring an amount of resource, a , from the other agent to the requesting agent according to the following formula:

$$a = \begin{cases} 0 & \text{if } drive\ level \geq threshold \\ 0.5 \times (threshold - drive\ level) & \text{if } drive\ level < threshold \end{cases}$$

where the threshold and level mentioned are those defined in chapter 5. Sharing policies

similar to this have been studied by Turner and Kazakov (2003) and have been found to be appropriate to this situation.

6.3.1 Resource Spacing

To assess the role that the availability of resources plays in determining the relative merits of communicative behaviour, the resource spacing experiment performed in section 6 is repeated with populations implementing the two new control behaviours described above. The first of these behaviours is identical to the standard behaviour, but communication between agents is not permitted. As this population is the only one not to act altruistically, it will henceforth be referred to as the *selfish population* for the purposes of brevity. The second control population is identical to the selfish population, but can share resources amongst themselves as outlined above. This population will be known as the *non-linguistic altruistic population*. The population implementing the behaviour described in chapter 5 will be called the *communicating population*.

From these experiments three 3-dimensional graphs are derived. In each case, the x and y axes are the distances d_{diff} and d_{same} respectively, while the z axis records the value of *relative benefit*_{AB} for the appropriate behaviours *A* and *B*. The first graph (figure 28) records the relative benefit of communication over selfish behaviour. Examination of this graph allows the identification of those environmental conditions under which language use is beneficial or detrimental. The second graph (figure 29) examines the relative benefit of communicative behaviour versus non-linguistic altruism. Through comparing these two populations, the effect of the altruistic aspect of the agents communication can to some extent be controlled for. This allows the benefits of linguistic altruism over a simpler form of altruism to be examined. Finally, figure 30 shows relative benefit of the non-linguistic altruistic population over the selfish population in order to evaluate the extent to which altruism can be seen as explaining the difference between communication and selfish behaviour, and therefore the extent to which language is implicated.³¹

³¹Please note that the graphs shown in figures 28–30 are rotated through 180° from the graph shown in figure 24; the rear corner of the earlier graph is equivalent to the front corner of the current ones.

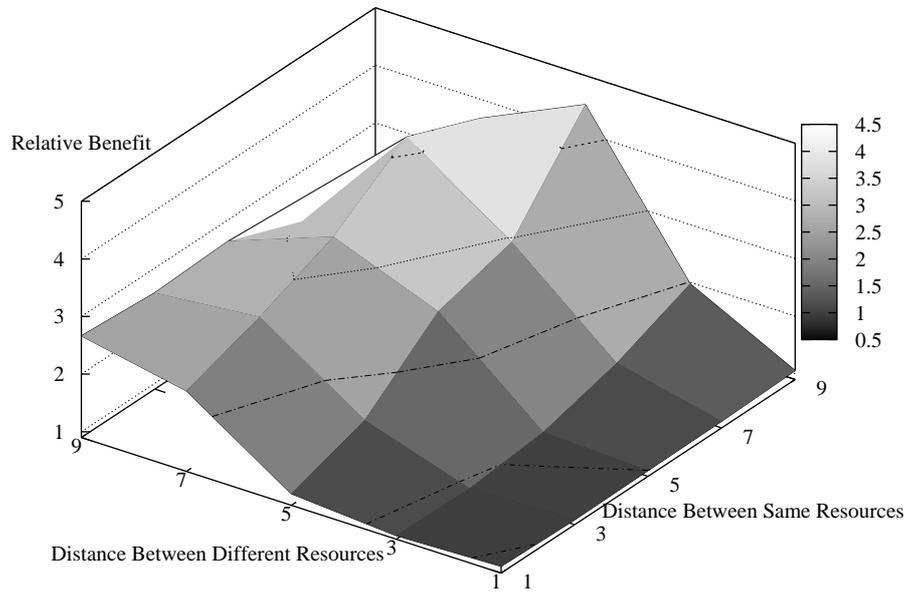


Figure 28: The relative benefit of language use over selfish behaviour as spacing varies. All results are averaged over 30 runs and 3000 time steps.

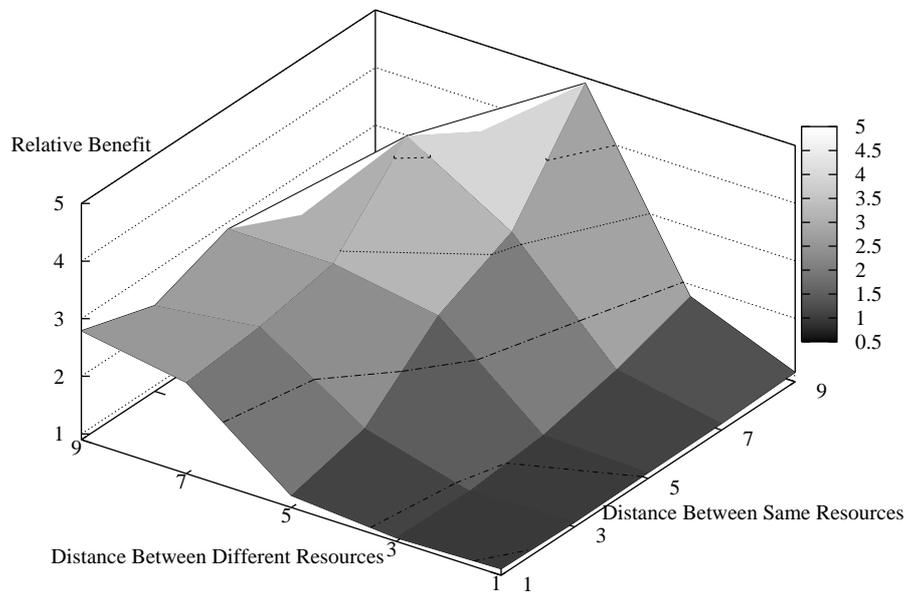


Figure 29: The relative benefit of language use over non-linguistic altruism as spacing varies. All results are averaged over 30 runs and 3000 time steps.

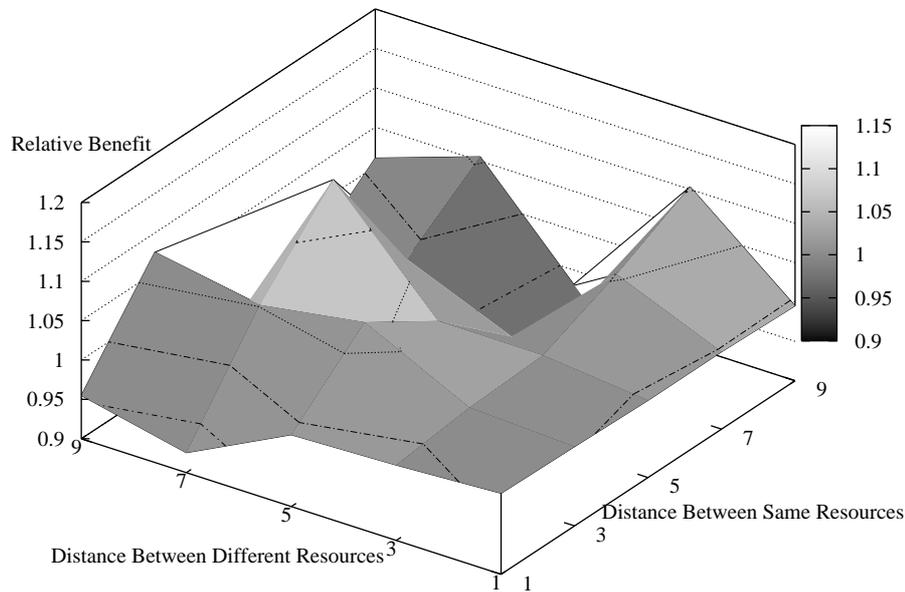


Figure 30: The relative benefit of non-linguistic altruism over selfish behaviour as spacing varies. All results are averaged over 30 runs and 3000 time steps.

The first and most important point to note is that the population able to utilise language always performs at least as well as both the selfish and non-linguistic altruistic populations. Additionally, this population performs significantly better than either of the other populations in a particular region of the studied parameter space. This region is not however that in which language using populations were found to grow to the largest size in section 6.2.

When the separation between food and water resources is at its least, the language using population was found to obtain its greatest size. However, at this point the relative benefit of language use over both of the other behaviours is virtually non-existent. This indicates that while the language using population was able to grow to a large size in this environment, language use was an unimportant factor in doing so. In contrast, the communicating population only reaches a modest size for $d_{same} = 9$, $d_{diff} = 5$ (around half of its greatest size), but has its greatest relative benefit at this point.

Language use also appears to offer benefits when d_{diff} is very large. However, figure 24

implies that language using populations head towards extinction in this situation. The apparent benefit of language use in this case comes from the use of average sizes to record population success rather than final size. All populations become extinct in this situation; however, the language using population does so more slowly. This results in a higher average size. If the simulation was to be run for a significantly longer period of time, the apparent benefit of language use would diminish for such situations as the average population size tended to zero for all behaviours.

In contrast to the language using population, which sees clear benefit for a certain region of the parameter space, there are no notable, significant differences between the two non-communicating populations. Figure 30 shows neither policy outperforms the other by a large amount at any point in the parameter space. The lack of any large deviation from a relative benefit of 1 in any situation, along with the lack of a coherent pattern, hints that these two behaviours perform equally, with the differences between them being mainly due to statistical noise.

The temporal behaviour of populations can be examined at points in each of the distinct regions identified; closely spaced resources, distantly spaced resources and in the region with the large relative benefit for language use. These graphs are shown in figure 31.

For closely placed resources, the temporal behaviour exhibited by non-communicating populations is virtually identical to that seen before for the communicating population. While the communicating population manages to reach a slightly higher peak size, the stable population size reached in the long term is equal for all population types. Similarly, when d_{diff} and d_{same} are large, the communicating population initially does slightly better, but all populations size converge eventually (to zero in this case). While figure 28 shows the behaviour of the communicating population to be relatively beneficial in this situation, this result is an artefact of the metric used, and would disappear if data were collected over a longer period.

The final graph of figure 31 is taken from the point at which the use of language is most relatively beneficial over both of the other behaviours. In this situation, the linguistic population converges to a stable long term population size of a little over 30

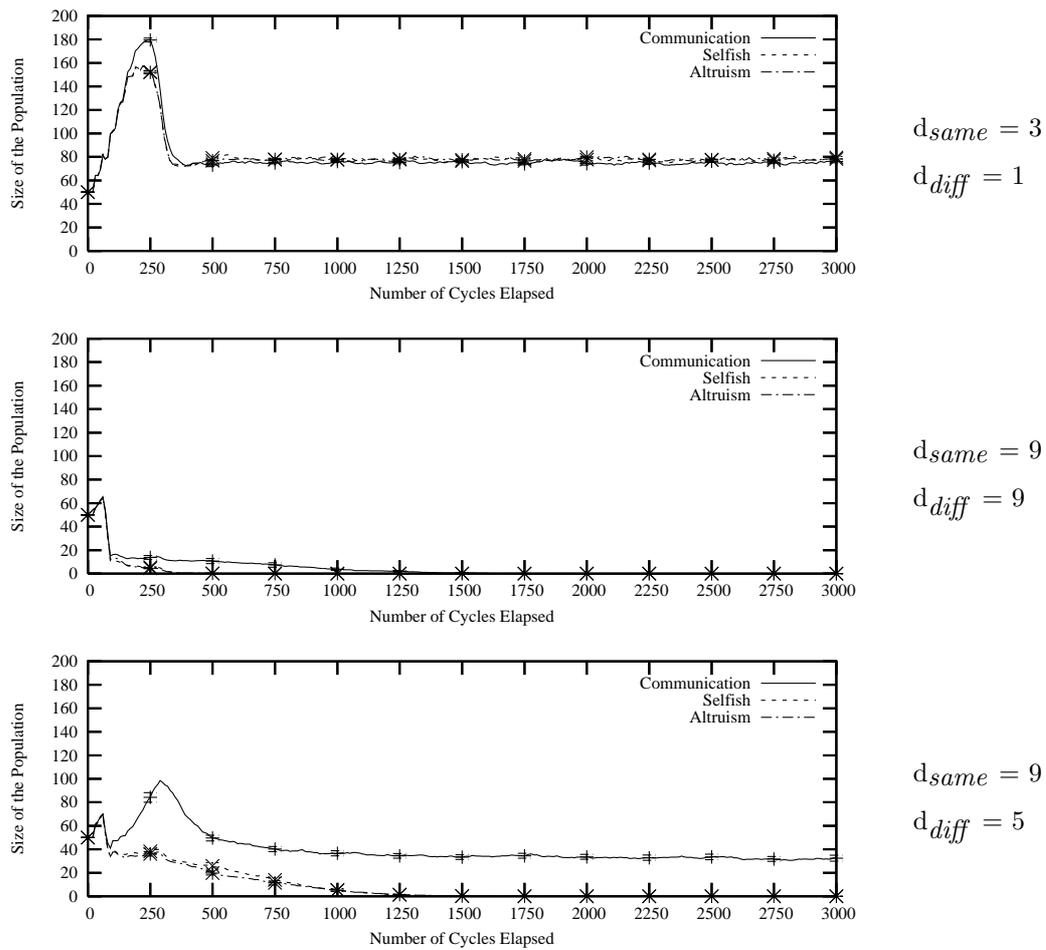


Figure 31: Comparison of populations implementing various behaviour types over time. From top to bottom, the graphs show results for points at which resources are closely spaced, widely spaced, and at which communication is most relatively beneficial. All results are averaged over 30 runs.

individuals. In contrast, the other populations fail to establish themselves and become extinct within the first 1500 time steps. In this situation, the use of communication not only leads to a larger final population size, but also makes the difference between the survival of the population and extinction. This suggests a strong source of evolutionary pressure for the development of communication in this setting.

6.3.2 Resource Volatility

As with the resource spacing experiments, the comparative benefits associated with language use in volatile environments can be found by supplementing the earlier results obtained for a communicating population with comparable results for selfish populations and non-linguistic altruistic populations. As above, this leads to three graphs comparing the relative benefits of these behaviour types.

The results previously shown for the communicating behaviour in volatile environments demonstrated interesting temporal behaviour. Therefore, results are also presented showing the temporal behaviour of all three different populations at given points in the two-dimensional space defined by the factors affecting volatility. This allows the impact on the population that occurs when resources move to be seen for each behaviour type.

The graphs of relative benefits obtained are shown in figures 32, 33 and 34. In each case, x and y axes show the initial level of each food resource and the time taken for a resource to decay completely from its initial state if not utilised faster. The z axes show the relative benefit of one policy versus another as indicated in the captions.

The most important results are those shown in figure 32. This graph reveals that the use of communication offers a relative benefit over not doing for all situations considered. Moreover, the use of communication proves massively beneficial in certain areas of the parameter space. When the parameter are set such that food is highly volatile or static, the use of language seems to offer negligible benefit. However, for a particular region where resources have display a reasonable amount of volatility, moving a few times on average per agent's lifetime, populations utilising language far exceed the

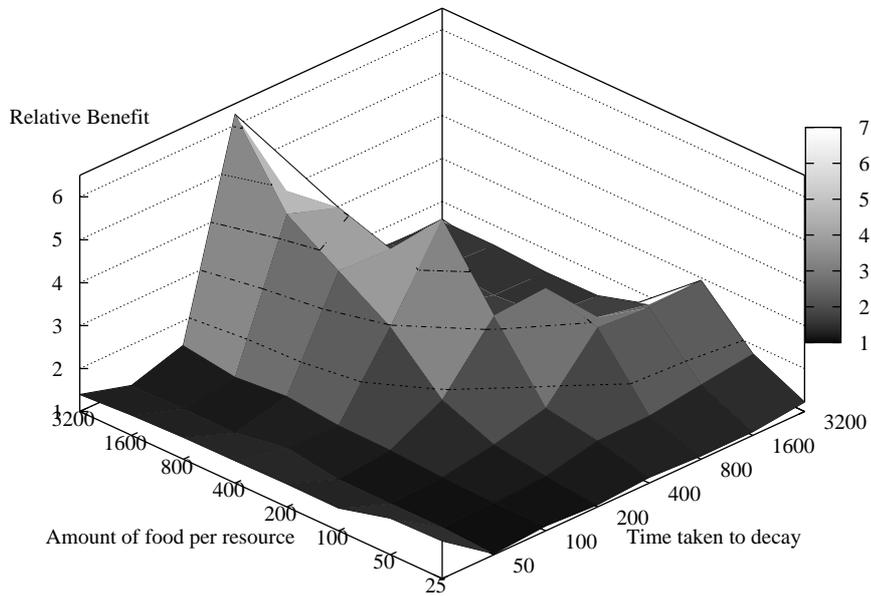


Figure 32: The relative benefit of language use over selfish behaviour in a volatile environment. All results are averaged over 30 runs and 3000 time steps.

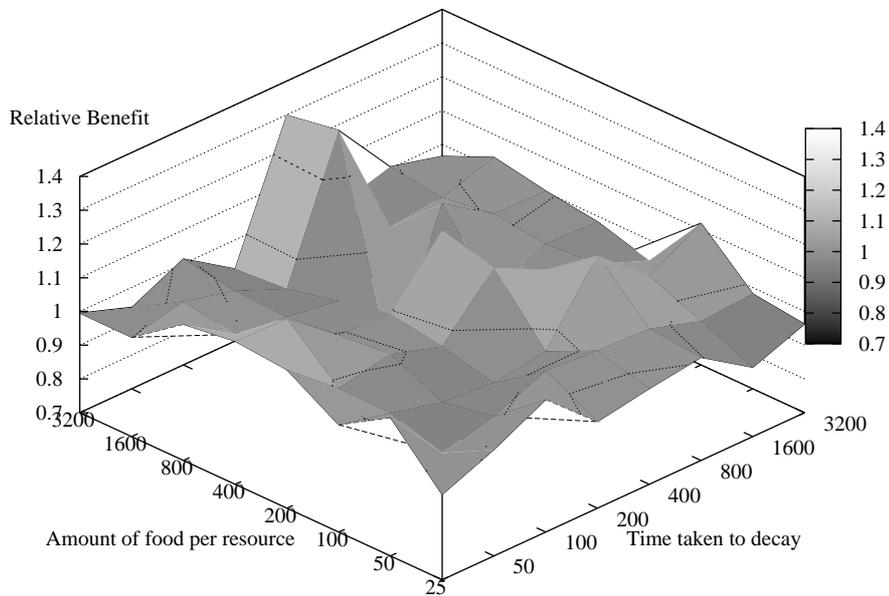


Figure 33: The relative benefit of non-linguistic altruism over selfish behaviour in a volatile environment. All results are averaged over 30 runs and 3000 time steps.

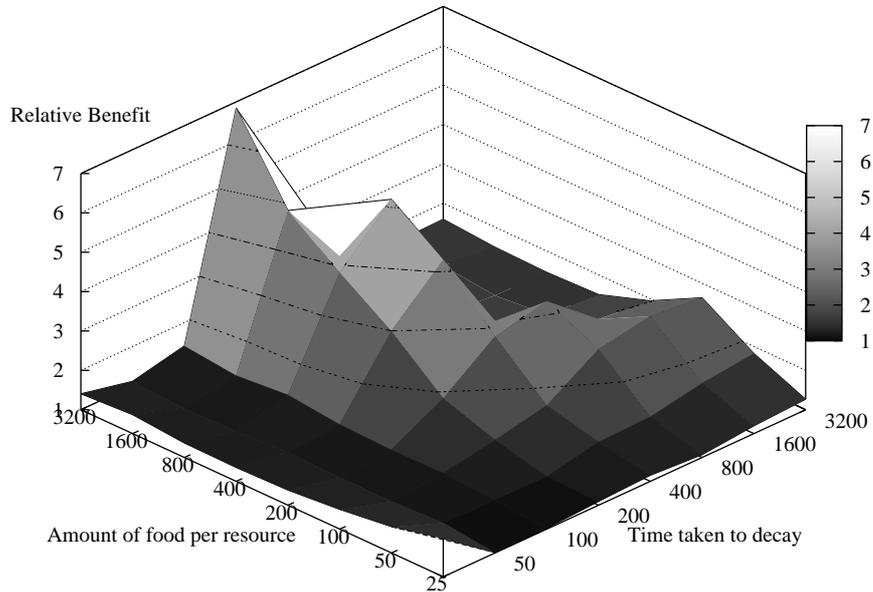


Figure 34: The relative benefit of language use over non-linguistic altruism in a volatile environment. All results are averaged over 30 runs and 3000 time steps.

performance of selfish ones; at the highest peak seen in the graph, the relative benefit of communication over selfish behaviour is 6.25. The largest relative benefit is seen in region where decay is reasonably faster but resource levels are high. However, a definite arc where language is beneficial can be seen through the environment from this peak area through to the zone where food is slow to decay and resource levels are low. This region is not the same as that in which language use was found to have the greatest absolute advantage. Comparing figures 26 and 32, it can be seen that while linguistic populations get larger the more static resources are, the greatest relative benefit is not derived in this situation. Rather, it is in a region of the parameter space where communicating populations only reach average sizes that their greatest relative benefit is manifested.

Examination of figures 33 and 34 confirms the massive benefit that language seems to offer over the other policies. Figure 33 fails to reveal any significant trend for the relationship between the environmental factors studied and the relative merits of selfish behaviour or non-linguistic altruism. There is some hint of a benefit for altruism over

selfish behaviour in the area of the graph near to where language use offered the greatest advantage. However, this benefit is on a much smaller scale than that seen for language use, and at the point at which language was most beneficial, the altruistic populations actually fare worse than the selfish one. Figure 34 confirms that any benefits due to language use far outweigh those gained through using the alternative non-linguistic form of altruism.

While examination of the temporal dynamics of the population sizes at each studied point of the x-y plane can be performed, presentation and analysis of this large amount of data would be very space consuming. Instead this temporal behaviour can be reasonably well shown through the presentation of graphs taken from each of the distinctly identifiable regions of the parameter space (figure 35). These regions are;

1. The region where food is highly volatile, due to either high decay rates or low resource levels. This region is seen across the front edges of the graphs in figures 32–34.
2. The region where food is highly static, because of low decay rates and medium to high resource levels. This region is seen at the right and rear of the graphs in figures 32–34.
3. The region where language proves most beneficial, i.e. that area where the decay period is around 200 timesteps and the resource levels are medium to high. This region is seen at the middle left of the graphs in figures 32–34.

Additionally, a graph is presented from the point adjacent to the last of these areas, at which the language using population reaches a sustainable long-term average, but the other populations die out.

The first graph of figure 35 shows what is typical behaviour for a point with very volatile resources. It was previously seen in section 6.2 that communicating populations in such conditions have a small initial increase in population size followed by a gradual but inevitable decline to extinction. This pattern is also followed by the other populations

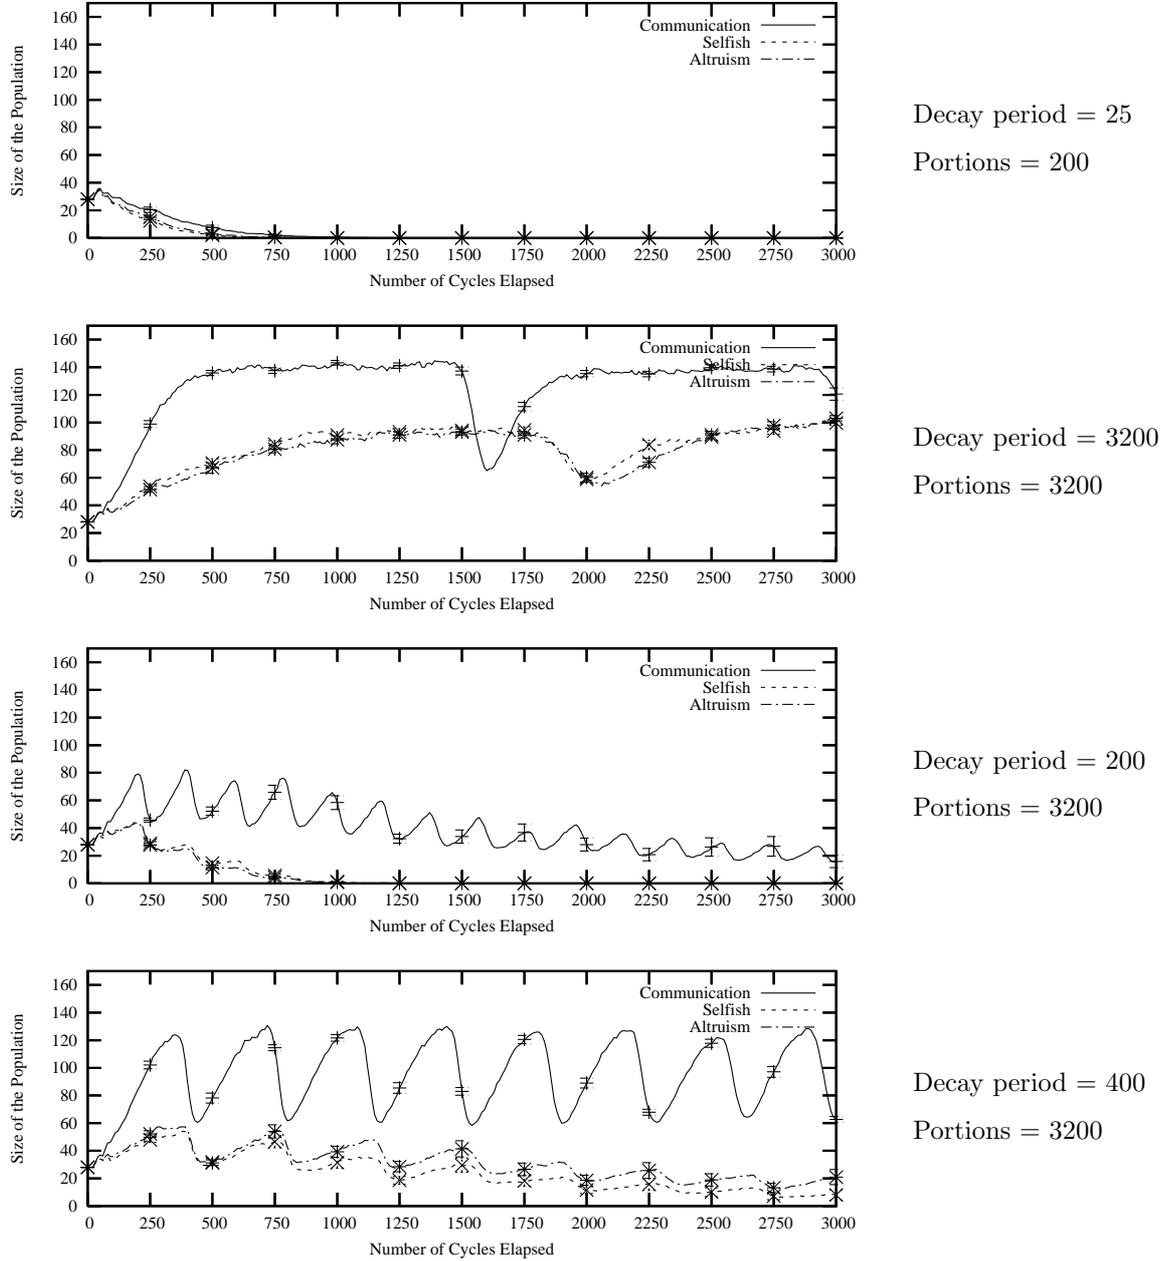


Figure 35: Comparison of populations implementing various behaviour types over time. From top to bottom, the graphs show results for points at which resources are highly volatile, least volatile, at which communication is most relatively beneficial, and in a situation in which communication makes the difference between survival and extinction. All results are averaged over 30 runs.

studied. Clearly in this region, the use of language is unable to offer any benefits over other behaviours.

The second graph is also typical of the region from which it is taken. In this region, all behaviours show a continual increase in population size, towards a maximum which depends on the behaviour implemented. Despite being relatively minor, the relative benefits brought through language in the case of highly static resources are liable to provide evolutionary pressure towards communication. Language use appears to increase the speed at which a population can grow, and also the size to which can be eventually reached. This rapid increase is seen from both the initial situation, and in the wake of the resource moving at around the midpoint of the run. In a situation in which a small population suddenly has the opportunity to expand dramatically such an ability could be promoted through evolution. Additionally, the faster rate of growth also appears to suggest that in the communicating population, more children survived to have children of their own, indicating individuals possessing a gene allowing linguistic ability could displace those lacking it from a population.

The most interesting graphs, from the perspective of the current work, are the final two from figure 35. These graphs are taken from the point at which communication provides the greatest relative benefit, and a point adjacent to this. From the first of these graphs, two things can be observed. Perhaps most obviously, the graph exhibits a periodic behaviour as mentioned earlier. While clearest for the communicating population, which reaches a near-stable average size, the periodic pattern is also overlaid on the generally declining population sizes of the other populations. Secondly, in non-communicating populations it can be observed that the number of agents falls slowly but steadily towards extinction, whereas the number of agents in the population using communication almost stabilises (except for the periodic increases and decreases due to resource volatility). In the final graph, the population does stabilise, while the other populations again head towards extinction. It is in this situation, in which communication may not just be a factor that can increase a population's fitness but might make the difference between life and death, that the greatest selective pressure on language use due to resource volatility may arise.

In all the graphs shown in figure 35, in only the last one do the non-linguistic altruistic population and the selfish population differ. In this case, the altruistic population performs slightly better, though given the long-term decline in both of these populations, this ‘improvement’ offers no real benefit.

6.4 Summary

This chapter has used the experimental framework of chapter 5 to experimentally explore the link between the environment and the performance of a population within that environment. Two approaches to this were taken. Firstly, the impact of various parameters on the size of a communicating population was assessed. Secondly, the role of the environment in determining relative benefit of language use was found through comparison with control populations who were either selfish or engaged in a non-linguistic form of altruism. Very little analysis and discussion of the meaning of the results obtained has been included in this chapter. Such discussion is presented in chapter 9, along with the meaning of further experimental results from chapter 7.

One factor considered was the volatility of resources. It was shown that as resources became more static, the average population size increased, regardless of whether this decrease in volatility was due to large resource stores or slower decay. When compared to other behaviour types, distinct patterns emerged for three areas of the parameter space;

- For highly volatile resources, all populations became extinct, but communicating populations at a slower rate.
- For moderately volatile resources, the relative benefit of communication to a population was at its highest, including, at some points, survival of communicating populations when other populations became extinct. This effect was particularly pronounced when the major source of volatility was decay.
- For highly static resources, all populations survived and grew, though with com-

communicating populations increasing in size more rapidly and reaching higher final sizes.

Studies into the impact of the positioning of resources on the benefits of a behaviour were also presented. Two parameters were varied in these experiments; the distance between resources of the same type, d_{same} , and the distance between those of differing types, d_{diff} . Both factors proved to effect the survival prospects of agent communities, though the latter of these distances had a far more substantial influence. When language use was compared with the other behaviours, a zone again developed where language allowed a far superior behaviour. This region occurs for relatively large d_{diff} and is at its greatest where d_{same} is also large.

As sharing information through language in a population involves altruism by necessity, the effect of both environmental factors studied was also assessed on a population of agents acting altruistically through non-linguistic means. No notable differences were observed between the performance of these altruistic agents and selfish ones. These data indicate the advantages observed in the communicating populations are not be entirely due to altruism benefiting the population, but may rely on the specifics of sharing through the medium of language.

7 Comparison of Different Behaviour Types

Having studied the role of the environment in determining the benefits and relative benefits of language use in the previous chapter, the focus of this chapter is the question of how this communicative behaviour could have arisen. In section 6.3, communication was shown to have significant benefits over two simpler non-linguistic forms of behaviour in certain circumstances; this implies significant evolutionary pressure towards language use. However, in order to move directly from the simpler forms of behaviour explored to the communicative one, at least three changes to the individuals would be needed. Firstly, the ‘syntax’ processing component used for navigation would need to become available to the speech circuitry. Secondly, the intention to communicate directions would be required. Finally, a common lexicon would also be needed.

To address the problem of multiple changes being simultaneously needed to move from one behaviour to another, this chapter introduces a range of cognitive abilities which each build incrementally on the abilities available at the previous level. This permits a set of intermediate steps between the silent behaviour and the system of fully syntactic behaviour studied in the earlier sections. Each of these behaviours consists of a relatively small step-up from the previous one and thus may be considered evolutionarily plausible. Experiments are conducted with populations implementing each of these behaviours to illustrate how each behaviour may have offered an evolutionary advantage over simpler ones.

As a central tenet of this thesis is that the use of syntax in language is an exaptation of earlier navigational abilities, various navigational behaviours are explored. Each of these navigational behaviours utilises the same cognitive abilities as one of the levels of language introduced. While the focus of this thesis is the exaptation of syntax-handling abilities from navigational abilities, experiments involving these navigational behaviours can help to suggest whether it is first necessary to have a fully-developed, regular language equivalent navigation system before language becomes useful, or whether useful

linguistic abilities could be an exaptation of a simpler form of navigation.

The rest of this chapter will now present the experiments as follows. First, the navigational behaviours that agents may follow are presented in section 7.1. Following this, the analogous communicative behaviours are presented in section 7.2. Finally, the results of the experiments conducted with agents utilising these behaviours are presented in section 7.3.

7.1 Navigational Behaviour

In order to examine the possible evolutionary pathway to language, four different levels of navigational ability are developed. The behaviour followed by all agents is similar to that explained in chapter 5. However differences exist, primarily due to the different information available to each agent for use in locating resources. The first behaviour (type-1 navigation) involves recording no knowledge of the location of resources, and having to locate them through exploration whenever needed. The second (type-2) is to remember the type of and approximate direction to the last passed resource that was located. Type-3 navigation extends type-2 to allow whole routes to be remembered, but not to allow them to be joined together (with the exception of joining the short term memory to a rule). The fourth, and most complex, (type-4) is to use the method of collecting knowledge and planning routes described in figure 16 (on page 104) and utilised in previous sections. These navigational behaviours, which are explained in more detail below, are summarised in table 4.

The most basic of agents (type-1) follow a restricted form of behaviour shown in figure 36. The major differences between this behaviour and that implemented by agents in earlier sections arises from the fact that these agents do not have the ability to remember resource locations. Consequently, it is impossible for them to plan routes (as they have no knowledge to use for this task) or to ask for directions (due to the fact that they lack the cognitive ability to memorise routes). As a result, when hungry or thirsty, these agents have no option but to randomly wander if they are unable to see

Navigation Type	Description
Type-1	Agent remembers no resources. Random exploration used to find resources when needed.
Type-2	Agent remembers single resource. Direction to resource remembered as last landmark passed.
Type-3	Agent remembers multiple resources. Directions to resources remembered as a collection of routes. Routes can not be composed.
Type-4	Agent remembers multiple resources. Directions to resources remembered as a collection of routes.

Table 4: Forms of navigation used in the model.

```

hungry := false;
thirsty := false;
if hunger ≥ thirst and hunger > drive threshold then
  hungry := true;
else if thirst > hunger and thirst > drive threshold then
  thirsty := true;

if hungry = true then
  if food in current square then
    Eat food;
  else if food in adjacent square then
    Move into square with food;
  else
    Make a random move;
else if thirsty = true then
  if water in current square then
    Drink water;
  else if water in adjacent square then
    Move into square with water;
  else
    Make a random move;
else
  Make a random move;
  Attempt to mate;

```

Figure 36: The behaviour followed by type-1 agents.

resources in their immediate vicinity.

Two other minor differences also exist between type-1 and type-4 agents. Firstly, as agents are unable to remember the location of resources, they do not undertake a search for resources if they do not know of any, unlike type-4 agents (see figure 16). Secondly, if they are not seeking resources, type-4 agents only make random moves if they are currently occupying a square with a resource. This constraint does not apply to type-1 agents: type-1 agents always make a random move during their turn if they do not need food or water. Preliminary studies indicated that were type-4 agents to do this then they frequently accumulated long routes in their short-term memories which had to be retraced at the beginning of any planned route and hence severely impacted on their survival. As type-1 agents do not plan routes, their random movement does not result in the same problem.

The behaviour of type-1 agents is extended to type-2 agents by the addition of a very crude method of remembering the position of the last resource seen (figure 37). Type-2 agents are equipped with very limited memory which allows them to remember only the type of the last resource they viewed and the description of the last landmark that they passed. Using this information, agents are able to implement a primitive, and highly unreliable, method for returning to the memorised resource should they need to. Heading in the direction of the last landmark passed will tend to take the agent towards this resource. Once at the landmark, the agent makes random moves on each subsequent turn until a resource is found. As the random movement is biased towards travel in a straight line, the random moves made at this stage will have a tendency to be in the opposite direction to those moves made by the agent between initially locating the resource and arriving at the landmark for the first time, hence an agent acting in this way will tend to navigate in the approximate, if not exact, direction of the resource.

As type-2 agents have (limited) memory and processing capabilities, it becomes possible for them to communicate in order to inform others of what they know. The form of communications available to both these agents and other types will be explained below,

in section 7.2.

```

hungry := false;
thirsty := false;
if hunger ≥ thirst and hunger > drive threshold then
  hungry := true;
else if thirst > hunger and thirst > drive threshold then
  thirsty := true;

if hungry = true then
  if food in current square then
    Eat food;
  else if food in adjacent square then
    Move into square with food;
  else if last resource encountered was food then
    Head towards last landmark passed;
  else if can communicate and another agent is present then
    Ask for help;
  else
    Make a random move;
else if thirsty = true then
  if water in current square then
    Drink water;
  else if water in adjacent square then
    Move into square with water;
  else if last resource encountered was water then
    Head towards last landmark passed;
  else if can communicate and another agent is present then
    Ask for help;
  else
    Make a random move;
else
  Make a random move;
  Attempt to mate;

```

Figure 37: The behaviour followed by type-2 agents.

The ability to return to a previously passed resource can be further improved by enabling agents to retain lists of descriptions of all the landmarks that they passed as in the previous chapter. This is exactly what type-3 navigation allows agents to do. However, the full form of processing present in the previous chapter is not used, rather agents are unable to compose multiple rules together to form a single route. Such behaviour requires a greater amount of memory than type-2 behaviour; multiple landmarks descriptions must be remembered instead of just one. This also necessitates an increase

in processing abilities to allow basic list processing functionality, such as appending landmark descriptions and reversal of the stored route. No algorithm is provided for this behaviour type, as it is identical to that in the previous chapter (see figure 16 on page 104) except that the routes returned from the route planning algorithm are more limited.

The final type of navigational behaviour present in the system (type-4 navigation) is the full system of navigation used in the previous experiments (see figure 16 on page 104). As this behaviour was comprehensively explained in chapter 5, no further such description is repeated here. This behaviour can be viewed as building on type-3 behaviour to allow multiple route sections to be combined to form one longer route. Again, this requires an increase in processing power from the previous level of navigation; as shown in chapter 5, this form of navigation requires a cognitive capacity equivalent to that needed to manipulate a regular grammar.

7.2 Communicative Behaviour

As with the navigational abilities, four levels of communicative behaviour are also defined. Each of these levels utilises the same cognitive capabilities as one of the types of navigational behaviour and is equivalent to the form of communication that would be available if the abilities underlying that type of navigation were to be exapted by the language faculty.

The first behaviour (type-A) is to not carry out any communicative activity. This behaviour can be carried out regardless of which of the navigational behaviours is being undertaken. However, the other types of communication are only available to agents with certain levels of navigational ability, due to the need for particular forms of knowledge and certain processing abilities. The second behaviour (type-B) is to communicate lexically by describing a single landmark in the approximate direction in which a resource may be found. Type-C communication allows the verbalisation of the an entire path that the agent has travelled along. The final type of communication (type-D) in-

Communication Type	Description
Type-A	Agent can not communicate at all.
Type-B	Agent can communicate a single landmark description, indicating direction.
Type-C	Agent can communicate a sequence of landmark descriptions, indicating a route. Agent cannot produce a new sequence, only articulate an existing one.
Type-D	Agent can communicate a sequence of landmark descriptions, indicating a route. Agent can produce a new sequence, using its knowledge as a regular grammar.

Table 5: Forms of communication used in the model.

volves using the full power of the syntactic language described above to describe routes. As with navigational behaviours, these communicative levels are summarised in a table (table 5).

Unlike type-A communication, which can be implemented by any agents, type-B communication requires at least the abilities underlying type-2 navigation. When type-2 agents communicate in this way, the speaking agent will articulate a description of the landmark it most recently passed, if the type of resource it most recently saw was the one that the other agent requires. The agent receiving directions in this manner will proceed to the mentioned landmark as if they had generated this direction from their own memory. Such communication can be described as *lexical*: the description communicated is simply a list of words and has no internal structure.

It is also possible for this communicative behaviour to be implemented by agents with higher levels of navigation. In the case of such an agent being called upon to generate a direction, it will plan a full route, and then vocalise the first landmark in the route. When receiving this type of directions, an agent with route planning abilities will behave in the same manner as an agent with only type-2 navigational abilities until it reaches the landmark in question, when it will revert to acting at its full capacity.

Type-C communication is only possible between agents with type-3 or type-4 navigation. This form of communication has some aspects of syntax, in so much as it allows an

ordered list of landmarks to be translated into an ordered sequence of words, but lacks other aspects, specifically the ability to combine grammatical rules in speech. Agents communicating at this level can articulate a route which they have remembered, but cannot generate a new route combining multiple routes from their memory.

In order to communicate at type-D level, by using full syntax, the agent needs to possess the highest level of navigational ability. In the case of a speaker, this is necessary in order to be able to gather the knowledge and to plan the route to convey, while a listener requires this ability in order to correctly memorise the route and subsequently follow it. Type-D behaviour was that implemented by communicating agents in the previous chapter.

Given the type of navigation that an agent uses and its communication system, it is possible to categorise agents as a compound type. For example, an agent using type-A communication and type-4 navigation will be subsequently be referred to as a type-A4 agent for the purposes of brevity.

7.3 Experimental Results

In the previous two sections, a range of cognitive abilities which permit a selection of navigational and communicative behaviours were developed. These abilities lie along an axis between no capabilities and the ability to handle fully compositional regular language. The stated aim of this part of the thesis is to question how the grammar-handling abilities needed for modern human language could have arisen from earlier stages of human evolution at which they were absent; while the abilities outlined in the previous sections form a series of steps by which evolution may have proceeded (that is to say, each stage is relatively small and offers a presumably useful capability not present at the previous stage), no evidence has yet been presented that these cognitive abilities did evolve in this way. This section attempts to provide such evidence.

This evidence is achieved through performing experiments such as those used in chapter 6, but with populations implementing the behaviours outlined in sections 7.1 and 7.2.

Comparison of the resulting average population sizes can then be used to evaluate the relative worth of each of these behaviours and hence establish where evolutionary pressure exists to move from one such behaviour to another.

As explained above, not all combinations of linguistic and navigational ability are possible. This is due to the fact that certain navigational skills are needed to produce or understand the language used at each level. For this reason only 10 types of agent are available. Specifically, these are type-A1, type-A2, type-A3, type-A4, type-B2, type-B3, type-B4, type-C3, type-C4 and type-D4. The meaning of these behaviours is explained above in full and will not be reiterated here.

In chapter 6, *relative benefit* was defined and used to compare the performance of two populations. However, while this statistic provides a value summarising the relative merits of two behaviours, it is not entirely suitable for the required purpose here. As the values that the relative benefit is derived from are experimentally obtained, it may not be clear whether an apparent benefit of one behaviour over another is real or merely an artefact caused by the stochastic nature of the experimental results.³² For example, should one behaviour be found to be 5% better than another, should this be taken as evidence of a small improvement over the other behaviour or just as a statistical variation in the sample between two behaviours which are actually of equal quality?

In order to overcome this problem, rather than use the relative benefit to decide whether one behaviour outperforms another, a t-test is used to decide whether the means of two behaviours are statistically different. Through such a test, it is possible to find a *p-value* which, in this case, is the probability of that the observed values could be seen if the two behaviours truly have the same mean, based on the recorded sample means, variances and sizes.³³ P-values vary between 0 and 1, with the larger the value, the less inclined one should be to believe that the two behaviours really do have different

³²In the previous chapter, this was not so great an issue as it is here, as the plotting of a three dimensional surface allowed each result to be seen within the context of the parameter space as a whole and the broad pattern could easily be discerned.

³³The procedure for conducting this kind of hypothesis testing is well known, and can be found in Newbold (1996) amongst many other sources.

population means. In contrast, the smaller the value is, the more willing one should be to believe that the apparently better behaviour really is does offer an advantage. While no absolute p-value can ever state that two means are definitely different or the same, the values of 0.1, 0.05 or 0.01 are commonly used as thresholds for this purpose (the 10%, 5% and 1% significance levels respectively). When it is necessary to use a threshold in the current results, the value of 0.05 will be used.³⁴

For each pair of behaviours studied, it is possible to obtain p-values which reflect whether the apparently better of the two behaviours is likely to really outperform the other. However, as each level of ability has been designed to form an evolutionarily plausible step from the previous level, comparisons of the relative quality of each of these behaviours will focus on those separated by just one such step. While pairwise comparisons of the quality of each of these behavioural pairs can be conducted (and indeed are), most attention is given to those pairs of behaviours where a single step exists between them.³⁵ For example, in moving from type-A3 to type-B3, only a small change in the linguistic behaviour is needed, whereas in contrast, moving from type-A3 to type-C3 or type-B4 would require either a change of more than one step in one ability, or a change in both abilities respectively. It is possible to assess where evolutionary pressure exists to move from one behaviour to another (i.e. when one behaviour outperforms another behaviour); by focussing on those situations in which only a single step change occurs, it is possible to identify those situations where not only does this pressure exist but the change is so small that it may conceivably have occurred. Examination of these situations provides evidence on the evolutionary pathway that may have led to the ability to handle composition in language.

In addition to the p-values, which are related to whether one behaviour is better than another, it is also informative to present the improvement in one behaviour over another. These figures are calculated for each pair of adjacent behaviours and illustrate the

³⁴Far more detailed information on the use of significance testing and p-values can be found in the the book referred to in footnote 33.

³⁵For the sake of brevity, behaviours separated by a single step shall subsequently be referred to as *adjacent* behaviours. The reason for this term is obvious when the behaviours are shown in a graph or table.

relative improvement of the more complex of the behaviours over the simpler. This value is calculated as:³⁶

$$\text{relative improvement}_{AB} = \left(\frac{\text{average population size}_A}{\text{average population size}_B} - 1 \right) \times 100\%$$

While these figures cannot state whether a difference is significant, as the p-values do, they instead show the gradient across the space of implemented behaviours. When used in conjunction with the p-values this provides useful information. Should a difference between adjacent behaviours be significant, the percentage increase provides an indication of the relative size of this improvement, and hence the degree of evolutionary pressure to move from one to the other. The direction in which this pressure exists is shown by the sign of the percentage; as improvements are calculated for movement from the more simple behaviour to the more complex, a positive improvement implies a pressure towards the more complex behaviour, while for a negative sign, pressure exists towards the simpler.

In chapter 6 it was demonstrated that the environment in which experiments take place has a significant impact on the relative benefit of the behaviour used. Therefore, the experiments conducted in this chapter use a range of environments in order to ensure that any findings are not fortuitous artefacts of a single environment. Rather than use a large range of environments (as in the previous chapter where the aim was to assess the environmental role), a much smaller subset of environments is considered. As the impact of the environment has been studied in the previous chapter, it is possible at this point to choose a variety of environments for which language use has been found to have different relative benefits. This smaller set of environments is drawn from those used to study the effects of volatility in the previous chapter.³⁷ Four points in this parameter space were chosen for study and form a progression through this space from

³⁶These percentages are linearly related to the relative benefit measure used in the previous chapter. They are used here in preference to that earlier measure simply for the increased clarity which they bring to the presented figures.

³⁷The volatile environments are used in preference to the spacing environments as a clearer pattern of relative benefit was seen to emerge for these environments.

the most volatile region (where language use was found to have little if any benefit), to the point at which language use seen to be the most relatively beneficial in the previous chapter.

For the experiments in each of these four environments, three figures are presented which capture the behaviour of the different populations within that environment.

1. A table of average population sizes for each behaviour, along with p-values for adjacent behaviours.
2. A table of relative improvements for adjacent behaviours with those improvements which are statistically significant indicated.
3. A table of p-values for all pairwise comparisons of behaviour types, i.e. both those behaviours which are adjacent and those which are not.

The experiments presented in this chapter were conducted over a shorter number of timesteps to those in the previous chapter and hence the observed values are not directly comparable to the earlier experiment. The decision to shorten the runs was made in order to eliminate the long periods during which populations had become extinct, which occurred at the end of many of the results from chapter 6. This increases the importance of the time during which multiple agents are present in the environment and during which their communicative behaviour may occur. Specifically, experimental runs are conducted over 1000 timesteps, rather than the 3000 used in the previous chapter. It is clear from figures 31 and 35 that differences between populations implementing different behaviours have time to emerge during this shorter run time.

The first environment studied is taken from the most volatile part of the parameter space. Table 6 shows the average population sizes for each of the behaviours in this environment. Italicised values in this table show the p-values calculated for the pairs of behaviours which they are shown between. This table shows the general trend for more complicated behaviours to result in higher average population sizes, but is too complex to allow for the full appreciation of the significance of the data. Table 7 summarises

this data in an alternate format. Rather than show absolute values in this table, the relative improvement (which was defined above) is shown for all adjacent behaviours. In each case, this value indicates the improvement obtained in going from the simpler of the pair to the more complex. The values shown in bold indicate those differences which are statistically significant. The data can be comprehended more clearly in this table.

The data in this table shows significant differences between nearly all pairs of adjacent behaviours, though in some cases the simpler behaviour outperforms the more complex. The most important finding in these tables is that the benefit in acquiring a more complex navigational system almost always exceeds that of moving to a more complex communication system. It is only when the most complex navigation is in place that strong evolutionary pressure acts to increase the communication system. This suggests that rather than the two abilities evolving in step, navigation may have evolved first, followed by communication.

One particularly noteworthy result is the finding that in this environment, type-C communication outperforms type-D communication. By examining what abilities type-C communication gives, it is possible to suggest an explanation for this result. In essence, type-C communication allows the speaker to describe the route to the last resource that was visited. In volatile environments such as this one, this may be advantageous over the ability to discuss any resources which have been previously seen; by only discussing recently viewed resources, the agents are less likely to hear of a resource which has since moved, and hence waste time travelling to a no longer existing resource. This limitation on the abilities of type-C agents acts to suppress redundant information. While this type of communication is beneficial in an environment in which resources are volatile, it may be less useful in other environments, such as the one used to study spacing in chapter 6.

For the sake of completeness, a matrix of p-values for the difference between every pair of behaviours is presented in table 8. As would be expected from the previous 2 tables, the more complex behaviour almost always significantly outperforms the simpler.

7 COMPARISON OF DIFFERENT BEHAVIOUR TYPES

D						8.26
						<i>0.024</i>
C				4.21	<i>0.000</i>	14.09
				<i>0.002</i>		<i>0.010</i>
B		4.19	<i>0.004</i>	4.78	<i>0.000</i>	7.53
		<i>0.028</i>		<i>0.001</i>		<i>0.187</i>
A	4.06	<i>0.086</i>	3.8	<i>0.025</i>	4.17	<i>0.000</i>
	1		2		3	4

Table 6: Average sizes for populations implementing each behaviour in the first environment. Navigational behaviours are shown on the x axis and communicative behaviours on the y axis. Large values at the intersections of the behaviours show the average population size for agents implementing that behavioural combination. Small italicised values show p-values for the pair of behaviours they appear between.

D						⊗
↑						-41%
C				⊗	+234%	⊗
↑				-12%		+87%
B		⊗	+14%	⊗	+58%	⊗
↑		+10%		+15%		+8%
A	⊗	-6%	⊗	+10%	⊗	+67%
	1	→	2	→	3	→

Table 7: Relative improvement between populations implementing each adjacent behavioural pair in the first environment. Navigational behaviours are shown on the x axis and communicative behaviours on the y axis. ⊗s stand in the place of each implementable behaviour, with the values between them indicating the relative improvement of the more complex behaviour over the simpler. Values shown in bold are associated with p-values less than 0.05.

7 COMPARISON OF DIFFERENT BEHAVIOUR TYPES

	A1	A2	B2	A3	B3	C3	A4	B4	C4	D4
A1	-	0.086	0.447	0.492	0.000	0.315	0.000	0.000	0.000	0.000
A2	0.086	-	0.028	0.025	0.000	0.009	0.000	0.000	0.000	0.000
B2	0.447	0.028	-	0.901	0.004	0.908	0.000	0.000	0.000	0.000
A3	0.492	0.025	0.901	-	0.001	0.791	0.000	0.000	0.000	0.000
B3	0.000	0.000	0.004	0.001	-	0.002	0.000	0.000	0.000	0.000
C3	0.315	0.009	0.908	0.791	0.002	-	0.000	0.000	0.000	0.000
A4	0.000	0.000	0.000	0.000	0.000	0.000	-	0.187	0.005	0.022
B4	0.000	0.000	0.000	0.000	0.000	0.000	0.187	-	0.010	0.245
C4	0.000	0.000	0.000	0.000	0.000	0.000	0.005	0.010	-	0.024
D4	0.000	0.000	0.000	0.000	0.000	0.000	0.022	0.245	0.024	-

Table 8: P-values for comparisons of all behaviours in the first environment. Values indicate the likelihood that average of the behaviour in each column differs from the behaviour type in each row. Please note that as tests for inequality, rather than for one behaviour outperforming the other, were used, the table is symmetric along the leading diagonal.

D							12.16
							<i>0.000</i>
C					4.3	<i>0.000</i>	5.49
						<i>0.529</i>	<i>0.000</i>
B			4.14	<i>0.144</i>	4.44	<i>0.000</i>	11.49
			<i>0.089</i>		<i>0.002</i>		<i>0.010</i>
A	3.84	<i>0.746</i>	3.89	<i>0.597</i>	3.81	<i>0.000</i>	9.76
	1		2		3		4

Table 9: Average sizes for populations implementing each behaviour in the second environment. See figure 6 for the key.

7 COMPARISON OF DIFFERENT BEHAVIOUR TYPES

D						⊗	
↑						+122%	
C				⊗	+28%	⊗	
↑				-3%		-52%	
B		⊗	+7%	⊗	+159%	⊗	
↑		+6%		+17%		+18%	
A	⊗	+1%	⊗	-2%	⊗	+156%	⊗
	1	→	2	→	3	→	4

Table 10: Relative improvement between populations implementing each adjacent behavioural pair in the second environment. See figure 7 for the key.

	A1	A2	B2	A3	B3	C3	A4	B4	C4	D4
A1	-	0.746	0.031	0.826	0.004	0.005	0.000	0.000	0.000	0.000
A2	0.746	-	0.089	0.597	0.010	0.017	0.000	0.000	0.000	0.000
B2	0.031	0.089	-	0.018	0.144	0.327	0.000	0.000	0.000	0.000
A3	0.826	0.597	0.018	-	0.002	0.003	0.000	0.000	0.000	0.000
B3	0.004	0.010	0.144	0.002	-	0.529	0.000	0.000	0.000	0.000
C3	0.005	0.017	0.327	0.003	0.529	-	0.000	0.000	0.000	0.000
A4	0.000	0.000	0.000	0.000	0.000	0.000	-	0.010	0.000	0.005
B4	0.000	0.000	0.000	0.000	0.000	0.000	0.010	-	0.000	0.470
C4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-	0.000
D4	0.000	0.000	0.000	0.000	0.000	0.000	0.005	0.470	0.000	-

Table 11: P-values for comparisons of all behaviours in the second environment. See figure 8 for the key.

Experiments conducted in a slightly less volatile environment, show a similar pattern (tables 9, 10 and 11), though some behavioural differences which were found to be significant in the previous environment are no longer significant for this one. Again, the benefits of moving to more complex navigation seem to outweigh those of increasing communication complexity until the most complex navigational behaviour is reached.

Similar results also apply to the other two environments, for which syntactic language was found to be even more relatively beneficial (tables 12, 13 and 14), and maximally relatively beneficial (tables 15, 16 and 17).

No relationship can be observed between the relative benefit of syntactic language use (as determined in the previous chapter) and of the size of the improvements given by a particular behaviour. That is to say, for all environments, most of the observed differences between behaviours are significant, and the neither the size of the p-values nor the relative improvements appear to vary in line with increasing benefit of language use. From this it may be stated that the results appear to be relatively insensitive to the environmental conditions, and hence the finding that complexity in navigation appears to evolve before complexity in communication appears to be a general rule.

7.4 Summary

In chapter 6, it was seen that agents capable of communicating using a syntactic language were able to outperform those who implemented two forms of non-communicative behaviour. This chapter has built on those results to study the path by which this fully-syntactic communication may have evolved from silent behaviour. Additionally, the evolution of navigation to full route-planning capabilities was studied. In order to carry out this study, the computational model from chapter 5 was again utilised.

Within this framework, a range of communicative and navigational behaviours were defined, with each communicative behaviour having a navigational behaviour requiring analogous processing capabilities. These behaviours ranged from silence/wandering through to full syntax/route planning, and were carefully chosen such that moving

7 COMPARISON OF DIFFERENT BEHAVIOUR TYPES

D							50.96
							<i>0.000</i>
C				4.3	<i>0.000</i>	72.34	
					<i>0.272</i>		<i>0.000</i>
B		4.19	<i>0.081</i>	4.5	<i>0.000</i>	10.46	
			<i>0.091</i>		<i>0.013</i>		<i>0.529</i>
A	4.09	<i>0.348</i>	3.92	<i>0.350</i>	4.07	<i>0.000</i>	10.08
	1		2		3		4

Table 12: Average sizes for populations implementing each behaviour in the third environment. See figure 6 for the key.

D						⊗	
↑						-30%	
C				⊗	+1583%	⊗	
↑				-5%		+592%	
B		⊗	+7%	⊗	+132%	⊗	
↑		+7%		+11%		+4%	
A	⊗	-4%	⊗	+4%	⊗	+148%	
	1	→	2	→	3	→	4

Table 13: Relative improvement between populations implementing each adjacent behavioural pair in the third environment. See figure 7 for the key.

7 COMPARISON OF DIFFERENT BEHAVIOUR TYPES

	A1	A2	B2	A3	B3	C3	A4	B4	C4	D4
A1	-	0.348	0.489	0.910	0.028	0.209	0.000	0.000	0.000	0.000
A2	0.348	-	0.091	0.350	0.003	0.030	0.000	0.000	0.000	0.000
B2	0.489	0.091	-	0.357	0.081	0.500	0.000	0.000	0.000	0.000
A3	0.910	0.350	0.357	-	0.013	0.127	0.000	0.000	0.000	0.000
B3	0.028	0.003	0.081	0.013	-	0.272	0.000	0.000	0.000	0.000
C3	0.209	0.030	0.500	0.127	0.272	-	0.000	0.000	0.000	0.000
A4	0.000	0.000	0.000	0.000	0.000	0.000	-	0.529	0.000	0.000
B4	0.000	0.000	0.000	0.000	0.000	0.000	0.529	-	0.000	0.000
C4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-	0.000
D4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-

Table 14: P-values for comparisons of all behaviours in the third environment. See figure 8 for the key.

D						60.85
						<i>0.000</i>
C				4.07	<i>0.000</i>	79.43
					<i>0.014</i>	<i>0.000</i>
B			4.01	<i>0.006</i>	4.54	<i>0.000</i> 11.47
				<i>0.388</i>	<i>0.008</i>	<i>0.000</i>
A	4.03	<i>0.305</i>	3.86	<i>0.293</i>	4.03	<i>0.000</i> 9.14
	1		2		3	4

Table 15: Average sizes for populations implementing each behaviour in the fourth environment. See figure 6 for the key.

7 COMPARISON OF DIFFERENT BEHAVIOUR TYPES

D							⊗
↑							-23%
C					⊗	+1853%	⊗
↑					-10%		+593%
B		⊗	+13%	⊗	+152%		⊗
↑		+4%		+13%			+25%
A	⊗	-4%	⊗	+4%	⊗	+127%	⊗
	1	→	2	→	3	→	4

Table 16: Relative improvement between populations implementing each adjacent behavioural pair in the fourth environment. See figure 7 for the key.

	A1	A2	B2	A3	B3	C3	A4	B4	C4	D4
A1	-	0.305	0.895	0.960	0.006	0.795	0.000	0.000	0.000	0.000
A2	0.305	-	0.388	0.293	0.001	0.224	0.000	0.000	0.000	0.000
B2	0.895	0.388	-	0.860	0.006	0.710	0.000	0.000	0.000	0.000
A3	0.960	0.293	0.860	-	0.008	0.836	0.000	0.000	0.000	0.000
B3	0.006	0.001	0.006	0.008	-	0.014	0.000	0.000	0.000	0.000
C3	0.795	0.224	0.710	0.836	0.014	-	0.000	0.000	0.000	0.000
A4	0.000	0.000	0.000	0.000	0.000	0.000	-	0.000	0.000	0.000
B4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-	0.000	0.000
C4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-	0.000
D4	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	-

Table 17: p-values for comparisons of all behaviours in the fourth environment. See figure 8 for the key.

from one to the next constitutes an evolutionarily plausible step. By comparing the performance of populations implementing behaviours separated by a single evolutionary step, it was possible to determine which behaviours offered benefits over others, and how large those advantages were. Intuitively, the greater the advantage of one behaviour over another, the more likely evolution would be to select for the better behaviour given the opportunity.

The obtained results revealed that, on the whole, the strongest evolutionary pressure was in the direction of increased complexity in navigation. Only once the most complex form of navigation had evolved were large benefits in increasing the complexity of communication observed. These findings were remarkably consistent across a range of environmental conditions. If the earlier theory that language and navigation share a neurological underpinning is true, this provides evidence that these abilities were more likely to have evolved first for navigation and then been exapted for language use, rather than the possible alternative theory that increases in complexity occurred simultaneously for both abilities.

One counter-intuitive finding of these results was that fully syntactic language was usually outperformed by a slightly simpler behaviour, in which only the most recently visited resource could be discussed. The most likely explanation of this is that as only recently visited resources were discussed by agents with the simpler behaviour, agents received more reliable information in the volatile experiments studied.

8 Examination of the Emerged Knowledge

The previous chapters introduced a computational model of language evolution and then used this to show how populations with more complex forms of cognitive ability fared better. This was true when this improved cognition was used for either navigational purposes or for communication. This chapter now turns to studying the navigational and linguistic knowledge possessed by agents in the earlier simulations in an attempt to ascertain *why* more complex abilities resulted in improved population survival.

In analysing this knowledge, one can adopt two viewpoints. Firstly, the information can be seen in terms of the agents knowledge of it. From this perspective, the language and navigational knowledge are encapsulated in the grammar internal to each agent. In the second interpretation, the knowledge is regarded in terms of how it is used. In this case, the idea of knowledge is more closely aligned with the routes produced for communication and navigation.

These two views of the knowledge possessed by agents corresponds closely with the Chomskyan idea of I- and E-language (Chomsky, 1986). In identifying knowledge with the internal grammar known by agents, a viewpoint similar to that of I-language is adopted. In contrast, an analogue of E-language is seen when the agents knowledge is viewed from the perspective of the routes generated.

8.1 Different Classes of Knowledge Utilisation

In chapter 7, experiments were performed with agents implementing various types of communicative and navigational behaviour as summarised in tables 4 and 5. Using the experiments in that chapter, it was possible to determine which types of behaviour outperformed which other behaviours and to use that information to determine a pos-

sible evolutionary route from no syntax processing abilities to a capability equivalent to regular language handling, which was used for both navigation and communication. This fulfilled the second aim of the thesis as stated in the chapter 1.

In this section, attention shifts from the now fulfilled aims outlined early in this thesis and examines the side issue of the form of the knowledge gained by these various types of agent. Specifically the aim is to identify the factors of the knowledge possessed by each agent type that explains why one type of agent outperforms another. As mentioned above two approaches to this are taken, one viewing language as an internal phenomenon and one viewing it as an external phenomenon.

8.1.1 Knowledge as an Internal Phenomenon

One way of viewing a language³⁸ is through the grammar used to generate it. Applying this perspective to the current work results in appraising the level of agents' knowledge through examination of the internal grammars that they possess.

In doing so however, it is impossible to examine the grammars individually and draw conclusions from that. Some sample grammars are shown in figure 38. Clearly simply reviewing by eye the several hundred such grammars which may result from any statistically significant experimental setting would be time-consuming and highly subjective.

Therefore, instead of studying the grammars themselves, statistics related to the grammars are instead be used. There are various properties of the grammar that could be studied, but two of the more meaningful and interesting properties are the number of rules contained the grammar and the length of these rules. Mapping these quantities into the agents' problem domain is approximately equivalent to determining the degree to which the agent may knowledgably move around the environment (or communicate such movements) and the directness of such travel.

Specifically in looking at the number of rules known to an agent, only those rules

³⁸Where language is used here in its mathematical sense.

containing landmarks will be considered. This reduces the rules counted to those of the form:

$$\textit{goto}(\textit{PosX}) \Rightarrow \textit{goto}(\textit{PosY}) L_1 L_2 L_3 \quad (31)$$

$$\textit{goto}(\textit{food}) \Rightarrow \textit{goto}(\textit{PosB}) W_1 W_2 W_3 W_4 \quad (32)$$

and excludes those of the form:

$$\textit{goto}(\textit{food}) \Rightarrow \textit{goto}(\textit{PosX}) \quad (33)$$

The former are included while the latter are excluded as the aim is to document the rules known for movement around the environment; the latter type of rule records a different property, namely the location of endpoints to which one may wish to move. Furthermore, noting that the first two rules above are also inherently different (the first being acquired through exploration and the second through navigation) the number of each of these rules is collected and reported separately.

Finally, rules are broken down into classes depending on the type of their endpoints. There are three such classes possible, those rules which record a route between two resources, those recording a route between two non-resource holding locations,³⁹ and those for which one end is a resource location and the other is not. In determining whether a rule's endpoints are resource locations, the non-terminals in the rules are matched against the rules of the form of equation 33 known to that agent. Therefore, it is more accurate to state the rules are classified according to the number of ends that the agent *believes* are resource locations. It should be noted that rules of the form of equation 32 may only have at least one end at a resource location, though possibly two, depending on the nature *PosB*.

The second property of note is the length of the grammars' rules. Length in this case is taken to mean the number of landmarks mentioned by the rule, i.e. the number of terminals in the rule. This corresponds closely to the length of path (or subpath) that each rule could be used to generate. For the sake of consistency, the rules are again

³⁹For the sake of brevity, the term *waypoint* will be used to denote a location of which the agent has knowledge but which does not contain a resource.

broken down into the same classes as for the number measurements, i.e. by the method by which the rule was gained and by the type of the endpoints. As before, rules of the form of equation 33 are excluded from analysis.

Having identified interesting properties of the grammar and determined a way to measure them, all that remains is to determine in which experimental setting to take the measurements. Having seen in the previous chapter that a remarkably similar pattern of relative merit between agent types exists for those environments studied, here the reporting of results is limited to just one environment for the sake of brevity. Specifically, the environment is used in which the decay period is 400 and 3200 portions of food are available per resource.

Within this setting, a population is initialised and allowed to run for a set number of cycles, followed by the grammars of all agents remaining alive at that time being recorded and analysed. This is repeated 30 times for each behaviour type studied to allow for meaningful conclusions to be drawn. Only those agents of navigation type 3 or 4 are used in these experiments, as none of those at lower levels use the same internal grammar for knowledge, making comparison with such agents impossible. Based on preliminary experiments, a duration of 500 timesteps is used to ensure all types of populations frequently have enough surviving agents to draw valid conclusions about their grammars.

The results of such an experiment are shown in tables 18 and 19, the former showing the number of rules known to agents and the latter the length of those rules.

The result most immediately apparent from table 18 is that agents with communicative abilities of type C or D know substantially more rules on average than those of types A and B. This effect is most pronounced when all routes are considered. However, even when only those rules found through exploration are compared, the trend remains.

This may in part explain the results observed in the previous chapter. The advantage observed in higher forms of communication in chapter 7 may be attributable to the fact that such agents know a greater number of routes and are therefore more able to cope

Agent Type	Discovered Routes				Told Routes			All Routes
	Zero	One	Two	Total	One	Two	Total	
Type-A4	0.427 (0.819)	3.31 (3.83)	2.77 (3.59)	6.51 (6.23)	N/A	N/A	N/A	6.51 (6.23)
Type-B4	0.418 (0.813)	3.78 (3.96)	3.05 (3.88)	7.25 (6.59)	N/A	N/A	N/A	7.25 (6.59)
Type-C4	0.970 (1.42)	5.68 (3.69)	2.35 (1.87)	9.01 (5.13)	2.15 (1.58)	0.378 (0.503)	2.53 (1.65)	11.5 (6.47)
Type-D4	0.692 (1.16)	4.79 (3.19)	2.28 (1.63)	7.76 (4.68)	3.43 (1.56)	0.0802 (0.365)	3.51 (1.55)	11.3 (6.06)
Type-A3	0.433 (0.824)	2.47 (3.49)	2.60 (3.82)	5.5 (6.37)	N/A	N/A	N/A	5.5 (6.37)
Type-B3	0.241 (0.652)	2.66 (3.45)	2.69 (3.61)	5.59 (6.47)	N/A	N/A	N/A	5.59 (6.47)
Type-C3	1.34 (2.46)	5.73 (4.44)	2.68 (1.91)	9.75 (6.69)	1.92 (1.92)	0.292 (0.546)	2.21 (2.22)	12.0 (8.70)

Table 18: The number of paths known by different agent types, categorised by the method of acquisition and by the number of end points containing a resource. Figures shown for each entry are the mean and standard deviation of the number of rules known per agent.

Agent Type	Discovered Routes				Told Routes			All Routes
	Zero	One	Two	All	One	Two	All	
Type-A4	1.69 (1.26)	4.04 (1.25)	6.01 (1.93)	4.57 (1.04)	N/A	N/A	N/A	4.57 (1.04)
Type-B4	1.58 (1.39)	3.93 (2.17)	5.51 (1.65)	4.54 (1.41)	N/A	N/A	N/A	4.54 (1.41)
Type-C4	2.58 (1.31)	2.40 (1.06)	3.60 (0.666)	2.72 (0.814)	2.47 (1.04)	3.68 (0.711)	2.74 (0.936)	2.74 (0.737)
Type-D4	2.58 (0.904)	2.13 (0.786)	3.48 (0.610)	2.57 (0.657)	2.27 (0.508)	2.84 (0.833)	2.28 (0.495)	2.48 (0.498)
Type-A3	2.15 (1.92)	3.79 (1.53)	6.24 (1.78)	4.65 (1.24)	N/A	N/A	N/A	4.65 (1.24)
Type-B3	1 (0)	3.97 (2.22)	5.91 (1.35)	4.76 (1.11)	N/A	N/A	N/A	4.76 (1.11)
Type-C3	1.79 (1.03)	2.29 (1.12)	3.98 (0.996)	2.77 (0.900)	2.61 (1.31)	3.77 (1.17)	2.75 (1.29)	2.78 (0.819)

Table 19: The length of paths known by different agent types, categorised by the method of acquisition and by the number of end points containing a resource. Figures shown for each entry are the mean and standard deviation of the average length of rules known by agents.

with the exhaustion of a single resource.

In terms of the type of discovered rule that is known by agents, a difference again exists between those of types C and D and those of A and B. The grammars of the former are skewed towards possessing a greater proportion of routes between waypoints or from one waypoint to another than the latter. This is the case even when types A3 and C3 are compared; in this case even though C3 knows a greater number of direct routes, this forms a lower proportion of its total rules found through exploration.

This result is likely largely an artefact of the simulation used, as type C and D agents will form waypoints when receiving routes from others hence biasing their grammars to contain more rules making mention of these points. Type A and B agents are only capable of acquiring waypoints if they discover that a location which contained a resource no longer does so.

Examining the differences between agents sharing a form of communication but using different navigation types reveals that for types A and B, the more advanced form of grammar results in the accumulation of more rules. The distribution of the type of rule is however very similar. In contrast, type-C agents acquire slightly fewer rules for the more advanced form of navigation. The distribution of these rules is also altered to include a greater percentage of those acquired through communication rather than discovered.

The greater number of rules possessed by type A4 and B4 agents over their type-3 counterparts may explain their greater success in the same way that this factor may explain the relative merit of more expressive communication types. However, the apparent advantage of type-C3 agents over D3 ones does not fit with the previously obtained results. One explanation for this is that the examination of grammar in this way does not reveal how the grammar is utilised. This point will be returned to in the next section.

Attention turns now to an examination of table 19 and to the results concerning the length of rules. Again, a broad trend is immediately apparent from the results, specif-

ically that the grammar rules known to type C and D agents are much shorter than those known to the other agents. In fact the rules known are around 40% shorter on average.

This adds further evidence as to why those agents with more advanced language were seen to perform better in the earlier experiments. Shorter routes would enable an agent to travel to a required resource much faster and allow an increased probability of reaching the target before dying and presumably also allow greater time for mating. However, it should be noted that shorter rules do not necessarily entail shorter routes. This again highlights the need to study the grammar in action in the following section.

It is possible to find an explanation for this trend in the emergent behaviour of the agent populations. As communicating agents are far more likely to acquire their first rules from communication, especially with their parents, the first routes they gain are statistically most likely to be based on those already widely followed by the population. As agents use an algorithm for generating routes that is biased towards the shortest known route, the routes generated for the newly created agents will be biased towards shorter routes. Hence a process analogous to evolution is occurring in the rule set, in which the likelihood of a route being passed on to the next generation is related to an implicit fitness based on its usefulness and length.⁴⁰ Over time, the routes remaining are likely to converge to the shortest known in the whole population. In contrast, non-communicating populations do not transmit knowledge from one generation to the next and hence this “evolutionary” process does not exist and route lengths remain constant over time.⁴¹

8.1.2 Knowledge as an External Phenomenon

In contrast to viewing knowledge as a static commodity which is synonymous with the agents’ grammars, another formulation is to associate the knowledge with the way that

⁴⁰In this process (which is analogous as opposed to exactly identical to biological evolution), routes are analogous to phenotypes and the rules used to generate them to genotypes.

⁴¹This explanation is closely related to the concept of mimetic evolution suggested by Dawkins (1976).

it is used. Using this approach, while the grammar is involved in the knowledge, it is the method of utilising the grammar and the project of this that is seen as most important.

Again there are many aspects that could be studied from this perspective, but there are two which appear particularly relevant. Firstly, the length of the path resulting from each production is recorded. Secondly, in order to study the use made of the grammar, the number of rules involved in producing each path can be recorded.

In analysing the grammar, the length of rules was measured as an approximation to the length of paths that would result from it. Here, this length can be noted directly. As with the earlier method, the length is taken to be the number of landmarks in the path. As noted in chapter 5, some transformation of the rule that results from the generation process may occur, mainly to remove duplicated landmarks in the route. The length of path that is reported in this section is the path which remains after this transformation. As this is the route that agents will eventually communicate or follow, it is a better measure of the way that knowledge is utilised than the alternative. For a similar reason, the path length also includes any landmarks added as a result of the contents of an agent's short term memory.

The second aspect recorded of an agent's knowledge is the number of rules that it uses in planning a route. This information is analogous to the number of rules known of various types in the previous section. In both cases, the aim is to study the ease with which an agent can generate a route from one location to another. An advantage of this method however, is that it relates directly to the parts of the grammar that are actively used by the agent. In contrast, the grammar itself may contain much information which is rarely or never necessary.

As there are two methods by which a Type-3 agent may generate a route, the method used to quantify the number of rules used in a path is constructed in order to allow the distinction between the cases to be seen. This same method is also used for Type-4 agents for consistency. Specifically, if an agent is able to reach a resource by simply backtracking along its short term memory, this is classified as utilising 0 rules, even

though this requires knowledge held in a rule of the form $goto(food) \Rightarrow goto(PosX)$ from the grammar. In contrast, when a rule from the grammar directing the agent from one location to another is also used, this is classified as a 1 rule route. For this reason, the measure constructed should strictly be referred to as the number of rules used which contain terminals. However, repeatedly referring to it as such would severely impact on the readability of sentences and hence throughout the method of just referring to the number of rules shall be used.

In addition to the data mentioned above, the number of unsuccessful attempts to generate a route can be recorded in a similar manner.

The experimental setting used is identical to when the grammar was the subject of study. However, as the data to record is now produced in a dynamic manner and is not available for study by inspection of the agents, a different method of collecting data is required. In this case, whenever a route is generated during the simulation, the number of rules used and the resulting path length is recorded. When the simulation is terminated, all path generations that occurred within the last 100 timesteps are taken as the final dataset. The window of 100 cycles was found to provide a sufficient quantity of data through preliminary experimentation and is short enough to approximate the behaviour of the system at the termination point. The experimental runs from which this data is collected are the same ones from which the earlier data was obtained, ensuring comparison between the two is valid.

As earlier, the lack of grammatical abilities in agents of type 1 and 2 means that these are incomparable to type 3 and 4 agents, hence they are not studied.

The simulation results with regard to the number of rules used per path are shown in figure 39. This figure shows the number of rules used to generate each path as a percentage of all successfully generated paths for that agent type. Normalising by this quantity allows the distribution to be compared between different populations which may have contained differing numbers of agents and in which different numbers of paths were generated. The data is plotted on 3 graphs, one for the paths generated for navigation by type-3 agents, a corresponding one for type-4 agents and one for the

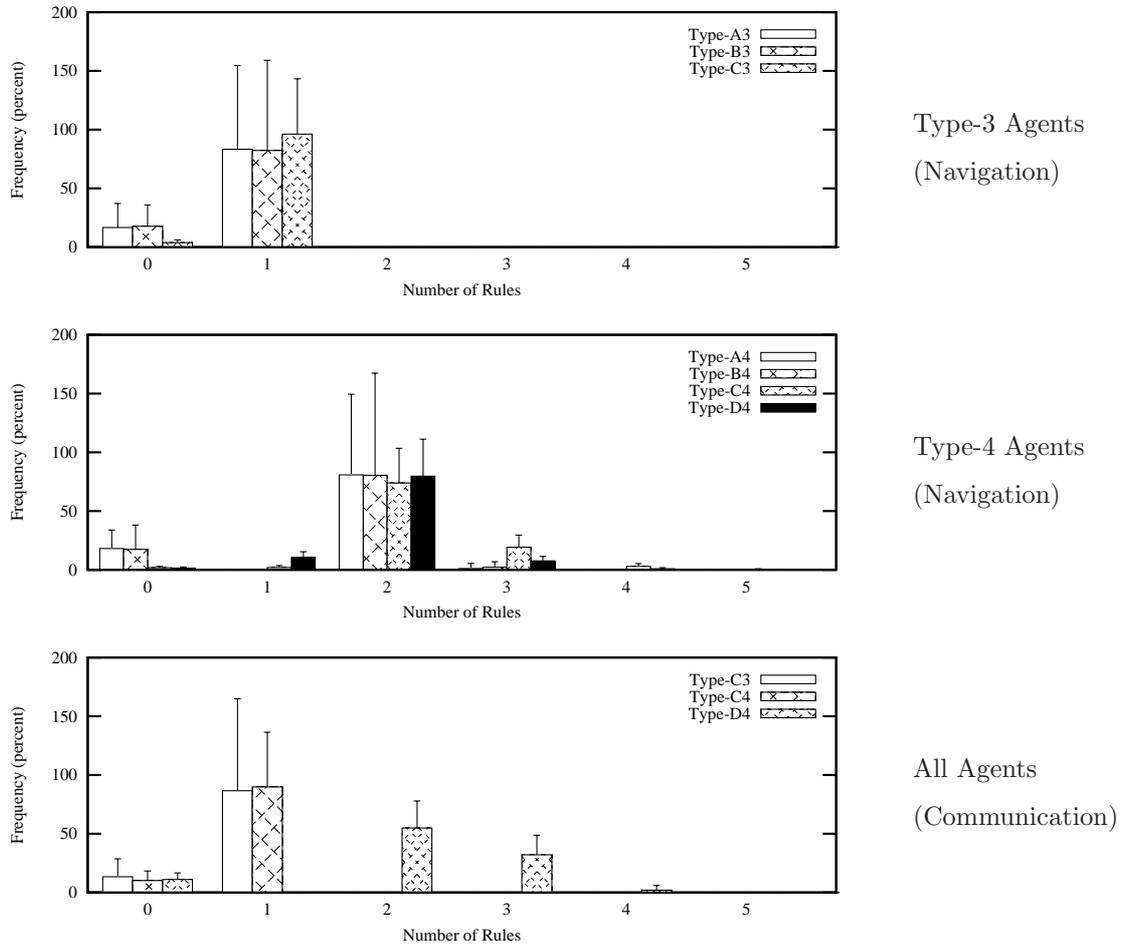


Figure 39: Distribution of number of rules per path used by agents. All numbers are averaged over 30 runs and are shown as a percentage of the number of paths successfully generated for that behaviour type.

communicated routes of all type C and D agents. Data is split in this way merely for presentational reasons⁴² and bars on different graphs can be directly compared to each other.

The most obvious result contained within this graph is the fact that type-3 agents use fewer rules per route than type-4 agents. However, this is an intuitive finding due to the fact that type-3 agents may never use more than one rule per route. Examining the results in greater detail reveals another trend towards type C and D agents using larger numbers of rules on average, and having almost exclusive use of routes of 3 rules or

⁴²Groups of 10 bars on the same axis being difficult to visually analyse.

more. Again this accords with an intuitive understanding of the issue. As these agents generate waypoints in more circumstances, it could be anticipated that they would need to generate routes from segments would lead not directly to or from resources, but between these waypoints.

The number of rules used to generate paths for communication follows a similar distribution to that used for generating them for navigation for both agents of type C3 and D4. For type-D3 agents though it is significantly different. This is due to the agent having differing abilities to generate routes for its own navigation and for communication to others. When navigating, the distribution is similar to type-D4 (with which it shares navigational abilities) while when communicating, behaviour similar to type-C3 (with which it shares communicative abilities) is observed.

These results help to clarify the results summarised in table 18 above. In the earlier table, the knowledge was viewed at a static level as defined by the grammar. The results here though highlight problems with that view. For example, the differing navigation and communication distributions exhibited by type-D3 agents would be very difficult to derive if the table was viewed in isolation. Likewise, not all rules possessed by type-C3 agents are usable; any rule between two waypoints will never be used in a path derivation, and other rules from waypoints to resources will only be of use if an agent can get to the waypoint without utilising other rules in the grammar.

The results for the lengths of generated paths are shown in figure 40. These are also normalised and subdivided in the same way as those of figure 39 for the same reasons.

The results show that, on average, the routes generated by type C and D agents have different distributions of length to those of type A and B. For the former two, a far higher percentage of routes of length one are generated. This may reflect the fact that when these agents have sufficient food or water, they remain in the vicinity of the last resource visited; the most probable place to find such agents is within a one landmark journey of a resource.

In contrast, type A and B agents generate a far greater proportion of their routes with

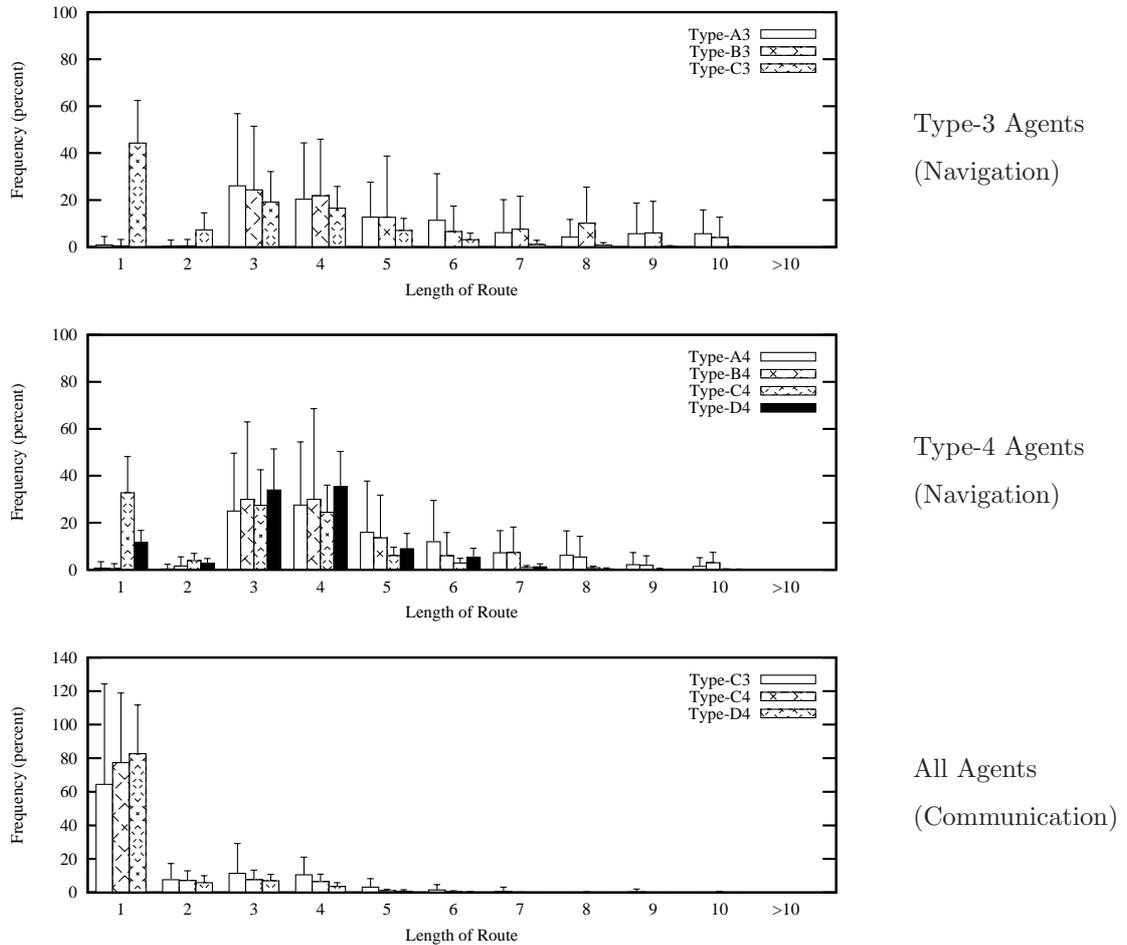


Figure 40: Distribution of length of paths generated by agents. All numbers are averaged over 30 runs and are shown as a percentage of the number of paths successfully generated for that behaviour type.

paths of length five or more. This accords with the results of rule length seen in table 19 and reinforces the conclusions drawn earlier: Agents with poorer communication skills perform worse as they follow longer routes leading to increased probability of death in transit and decreased time available for mating.

Similar comments can be applied to the communicated routes. Here, type-D agents outperform type-C3 agents by generating a greater percentage of routes of length one and fewer of higher length. However the difference between them is not as exaggerated as for the difference between types A and B, and types C and D. The high proportion of short rules can probably again be explained by the likely proximity of agents to

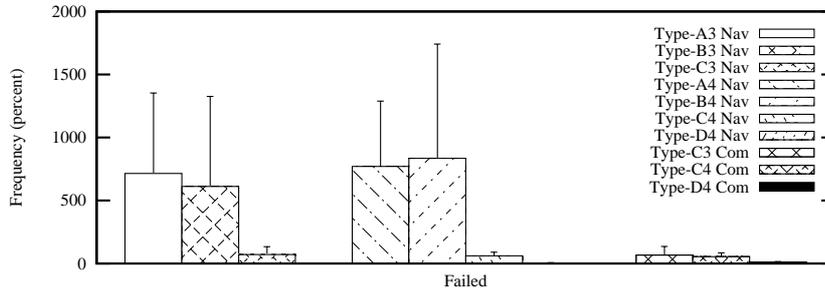


Figure 41: The number of path generation attempts that failed. All numbers are averaged over 30 runs and are shown as a percentage of the number of paths successfully generated for that behaviour type.

resources.

The final results for this experiment can be found in figure 41. This shows the number of attempted path generation attempts that were unsuccessful in creating a path. As with earlier figures, these are shown as a percentage of successful generation attempts to allow fairer comparison.

The figure reveals one final aspect that also explains the relative success of agents which can communicate paths. Type A and B agents fail to generate paths an overwhelming proportion of the time. In contrast, type C and D agents are able to generate paths in the majority of cases where they need them. An ordering exists whereby type-D4 agents are more successful in generating paths than type-D3 agents, who in turn are more successful than type-C3 agents. This holds true for both navigational and communicated routes, and mirrors the relative success of agent populations obtained in chapter 7.

8.2 Language as an Altruistic Act

Having investigated the knowledge associated with different agent types in chapter 7, this section now turns to looking at the knowledge present in the earlier setting of chapter 6. In order to minimise duplication of the results presented in section 8.1, this section will focus on a setting in which resources are non-volatile, but rather have a

given fixed spacing between them, such as those shown in figure 22 on page 122.

Specifically an environment is used which has parameter values of $d_{same} = 5$ and $d_{diff} = 5$. Other than the changed environment and the use of the agent types studied in chapter 6, the simulation setting remains the same. The same length of simulation run is performed and the same properties are recorded for both viewpoints of knowledge.

As it is unnecessary to reiterate the type of data collected and the reasons for doing so, the following sections proceed directly to presenting the experimental results.

8.2.1 Knowledge as an Internal Phenomenon

The analysis of the final grammars is presented in tables 20 and 21. As with the earlier results, the first table reports the number of rules known to each agent type and the latter records the average length of such rules.

Table 20 shows a distinction between the grammatical knowledge of communicating agents and non-communicating agents. Those agents who communicate know a greater number of routes, an average of around 50% more. However when linguistically acquired rules are removed this falls to a little over 10%. Additionally, the communicating agents' grammars feature more discovered routes not directly between resources. Little difference exists between the two types of non-communicating agents.

The results agree with both the findings relating to communicating agents earlier in this chapter and with the previous results from chapter 6. In section 8.1 the communicating agents were found to have a greater number of rules as here. This was also found to correlate with improved relative performance. Again this pattern is observed in this dataset. The increased number of paths leading to or from waypoints is again likely to be a result of higher levels of waypoint creation in the communicating agents.

The results shown in table 21 also mirror the findings from earlier in the section. In both cases, the communicating agents are observed to know shorter paths than their non-communicating counterparts. This effect is attenuated for routes directly between

Agent Type	Discovered Routes				Told Routes				All Routes
	Zero	One	Two	Total	One	Two	Total		
Communicating	0.420 (1.12)	2.70 (2.07)	1.52 (0.894)	4.64 (2.66)	1.90 (0.684)	0.153 (0.479)	2.06 (0.564)	6.69 (2.97)	
Silent	0.218 (0.767)	1.94 (0.998)	2.00 (1.42)	4.15 (2.04)	N/A	N/A	N/A	4.15 (2.04)	
Altruistic	0.189 (0.707)	1.93 (0.965)	1.99 (1.40)	4.11 (2.01)	N/A	N/A	N/A	4.11 (2.01)	

Table 20: The number of paths known by different agent types, categorised by the method of acquisition and by the number of end points containing a resource. Figures shown for each entry are the mean and standard deviation of the number of rules known per agent.

Agent Type	Discovered Routes				Told Routes			All Routes
	Zero	One	Two	All	One	Two	All	
Communicating	2.78 (1.52)	2.66 (1.49)	4.02 (0.914)	3.20 (1.08)	2.44 (0.869)	2.94 (1.31)	2.45 (0.834)	2.93 (0.846)
Silent	2.44 (1.68)	2.71 (1.99)	4.40 (1.31)	3.48 (1.35)	N/A	N/A	N/A	3.48 (1.35)
Altruistic	2.48 (1.74)	2.64 (1.95)	4.48 (1.31)	3.49 (1.40)	N/A	N/A	N/A	3.49 (1.40)

Table 21: The length of paths known by different agent types, categorised by the method of acquisition and by the number of end points containing a resource. Figures shown for each entry are the mean and standard deviation of the average length of rules known by agents.

one resource and another. In table 21 the difference in route length is much smaller.

One possible explanation for the decreased advantage associated with communicating agents in the present environment is that resources are more readily discovered in this setting. Recall the form of the layout (as typified by figure 21 on 121) and the forward biased movement followed by exploring agents. Given these two aspects in tandem, discovery of resources through wandering will often lead to near direct routes between them; any path from top to bottom of the environment (or vice versa) will allow resource discovery, and agents are biased towards taking such paths from edge to edge. In contrast, there is much greater potential to fail to discover resources when wandering in the environments of section 7.

8.2.2 Knowledge as an External Phenomenon

As there are only three types of agent in this setting, and it is only necessary to divide the paths produced by one of these agents into communicated and navigated ones, each set of results can be shown on its own single set of axes. Aside from this presentational issue, the information collected is the same as in section 8.1. This information is shown in figure 42.

The first graph of figure 42 shows the number of rules used in generating routes. As with the previous analysis, it is impossible to find any significant difference between the two non-communicative behaviours. The two aspects of the communicating agent are however different from this, and from each other. The distribution of the number of rules used by each of the agents for navigation is similar to those found for equivalent type-4 agents in the previous section. Similar comments therefore apply as to why this may be and what it reveals about the results in earlier chapters. However, the utterances produced by the communicating agents are very different and therefore must be examined in more detail.

By definition, in order to generate a communicated route using no rules, an agent must have just travelled from a location at which the sought resource can be found. Due

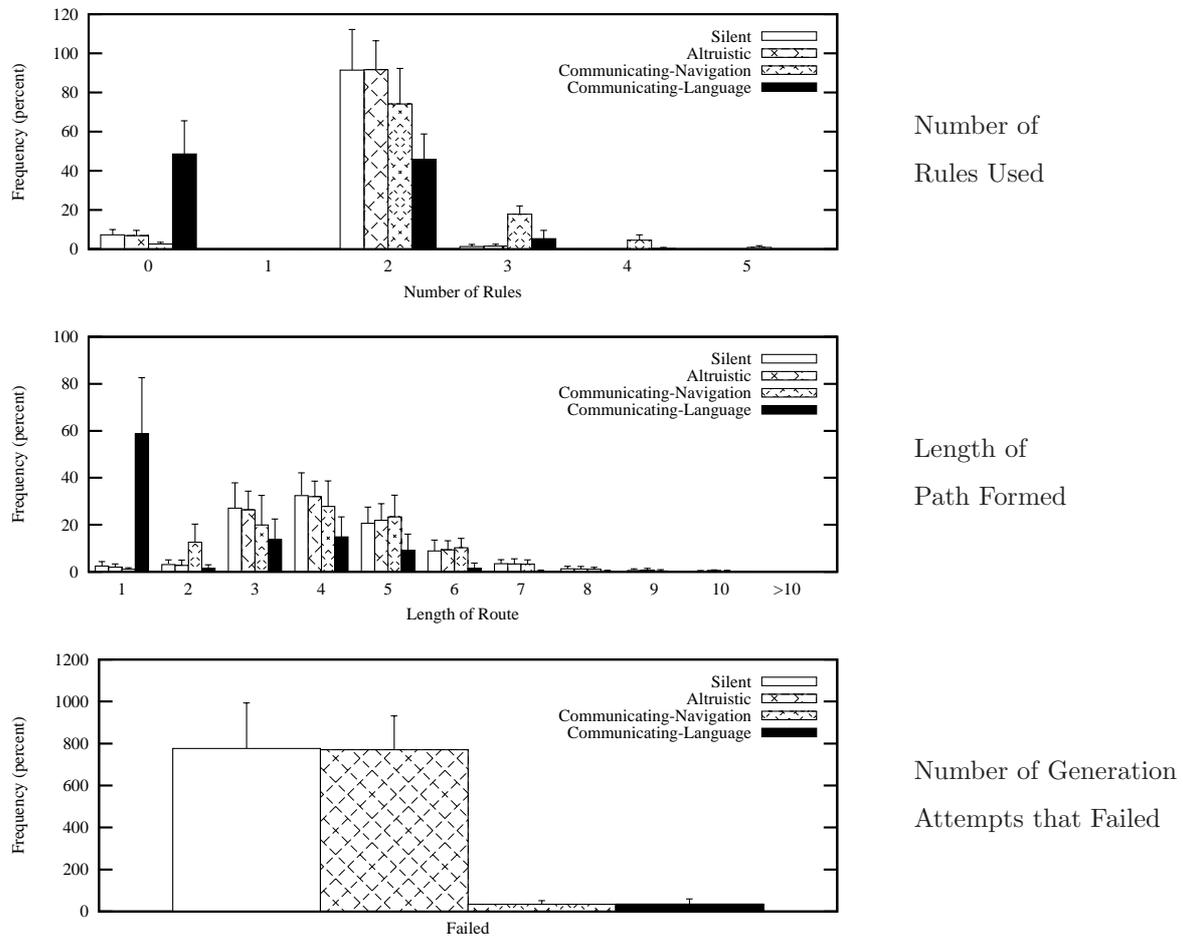


Figure 42: Distribution of number of rules used and length of paths produced by agents. All numbers are averaged over 30 runs and are shown as a percentage of the number of paths successfully generated for that behaviour type.

to the static nature of the resources in this environment, this will happen in around half of all requests. In contrast, in the earlier results in section 8.1, agents spend a disproportionate amount of time travelling to and searching for one of the two resources required (the volatile food). This means that an agent is much more likely to be asked to generate directions to food when its last visited resource was water. Hence this peak does not appear in the earlier figures.

Looking at the length of routes generated reveals each type has a similar distribution, with two exceptions. Firstly, a massively disproportionate amount of the communicated routes are of length one, though the remainder follow a similarly shaped distribution

to the other route types. Secondly, the navigational routes for communicating agents feature noticeably more routes of length two, with a slight decrease of routes of length three to compensate.

In the case of the first anomaly, an analogous pattern was noted in the utterances produced by those agents which communicated in the section 8.1 and the reasons given for that apply here too. In contrast, the other deviation has no direct analogy in the previous results. It is most likely caused by the communicative behaviour of this agent type. To travel from one resource to another in the studied given environment requires a path of length three or more. Therefore, any path of length one or two must have been generated within the central area of the environment. For non-communicating agents this is rarely done, possibly only in an initial wandering phase. However for communicating agents it may occur if they are told of a route in such a position. They will store the route, then subsequently generate a path based on this route in the following turn.

Finally, examining the proportion of route generation attempts that fail to produce a route is again instructive. The number of failures is an order of magnitude higher for non-communicating agents than for communicating ones. This strongly suggests that communicating agents have more knowledge about their environment, or are able to acquire such information easily from others when it is needed.

8.3 Summary

Having shown in previous chapters that the evolution of a compositional grammatical language from simpler types of representation was plausible, this chapter has focussed on the knowledge that is gained by agents in such a system. The aim in this chapter has not been to further strengthen the evidence gathered in previous chapters, but rather to seek to understand why the obtained results have been seen.

One view adopted in the chapter was to consider the knowledge possessed by agents as equivalent to the grammars that they contained. Specifically, the size of these gram-

mars was considered as well as average rule length. Such analysis found that those types of agent which had more complex communicative systems knew larger numbers of rules and that the rules they knew were on average shorter. Such findings hint at communicative agents having a greater knowledge of the environment and possessing more direct routes around it.

A second viewpoint was taken by not limiting the concept of knowledge to the grammar possessed, but to the way in which it is used. In this system, similar results were obtained, showing that the more communicative agents used a greater number of rules to generate each route, but that they generated fewer very long routes. It can again be concluded that communicating agents seemed to be following shorter paths around the environment, leaving them less prone to exhausting their resource levels and leaving longer for mating. The greater knowledge possessed by these agents was again evident, this time in the substantially fewer failed attempts at path generation that they suffered.

9 Conclusions

In this thesis, the aim has been to explore the possible origin of syntax in human language, and specifically one theory linking this feature to navigation. As was stated at the beginning of the introduction, syntax is one of the most important aspects distinguishing human language from other animal communication systems. Indeed, it is so fundamental to language that many researchers modelling language evolution have simply assumed the existence of syntax-handling capabilities without even explicitly mentioning this assumption.

Throughout this thesis, a pair of complementary aims have been pursued. This first of these has been to argue that the syntax-handling abilities used for language are an exaptation of those used for handling route planning in navigation. While the theory that navigation may be a precursor of some linguistic abilities is not entirely novel, the consideration of this idea in this thesis is more fully developed and more rigorous than any previous suggestions which have been made on the subject. The second aim has been to show how a mutation creating a link between this processing centre and the speech circuitry could have spread through a population due to the benefits it provides. These theories in combination offer an explanation of both how syntax in language arose in individuals and how it arose in the species.

The computational approach used is significantly different from many of the more recent simulations in the evolutionary linguistics field, and in some senses has more in common with the original work carried out in this field by Werner and Dyer (1992) and Ackley and Littman (1994). While the aim of this work is to study syntax (as opposed to signal coordination in these earlier works), the simulations embrace their core idea of studying language as a tool for some purpose, not as the end in its own right. This entails a situation in which the benefit of language use to a population can be measured through its utility in the tasks faced by a population and not through some externally defined concept of what a ‘good’ language is.

This avoids some of the problems that critics such as Bickerton (2007) have of the field,

specifically that in the more famous computer models, agents have to pass through a ‘babbling’ stage where no-one understands each other. The payoff for this behaviour can come only after much time or many generations, and it is not clear why individuals would persist with this behaviour. In contrast, the approach adopted here evaluates how language aids with the task of survival and reproduction; if language is useful, the benefit will be seen in improved performance at these tasks. If it is not, then that also will be observed.

A second difference between the approach used and the majority of current work is to focus on the language faculty and not the language. The prevailing trend in simulations in this field is towards modelling the coordination of lexicon or emergence of grammar in a population of agents who do not themselves evolve.⁴³ In these models, it is assumed that the language faculty pre-exists or that language learning uses domain-general learning mechanisms. While these models can illustrate the mechanisms that exist for self-organisation and coordination, they do little to illustrate how language actually began or why the agents in these models have evolved the cognitive capabilities that they have.

In contrast, this thesis has concerned itself with the origins of syntactic capabilities in language use. In doing so, the emphasis has been entirely on the evolution of (an aspect of) the language faculty, and the language used has been largely ignored. Indeed, the work is highly unusual in the field of language evolution in that the lexicon used by agents in a population is predefined and does not change throughout an individual’s or the population’s lifetime.⁴⁴

Having recapped the broad aims and approach of this thesis, this chapter now proceeds to present an overview of the major arguments and findings from earlier chapters. This occurs in 3 parts. Firstly, the conclusions drawn from the survey of the existing literature in chapters 1 to 4 are presented, with particular emphasis given to chapter 3.

⁴³Even in models in which agents do evolve, it is often simply a change in the choice between several equally complex behaviours, rather than between behaviours requiring different cognitive capabilities.

⁴⁴The grammar in some sense can be said to change, as some sentences become ‘grammatical’ as agents explore the environment and discover routes which may be communicated.

Next, the important novelties in the design of the computer model are shown, followed by a section in which the results obtained using this experimental setting are restated and their significance discussed. The chapter concludes with an outline of future work which builds on the theories and model developed in this thesis.

9.1 Detailed Summary of the Thesis

9.1.1 Literature Review

A review of previous work in four areas was presented in earlier chapters. Current theories on the nature and evolution of the linguistic faculty were presented in the introductory chapter. This was followed by an exploration of research into navigational behaviour, with a strong emphasis on animal-based experimental studies. Chapter 3 presented evidence on the grounding of language in the brain, focussing on research which demonstrated a link between language processing abilities and other non-linguistic capabilities. Finally, an overview of important techniques and results in the field of computer modelling of language evolution was presented.

The most important novel insights from this background material came in chapter 3. By linking together research in this chapter which is not traditionally part of the same discipline, it was possible to convincingly argue that linguistic and certain non-linguistic processing takes place using the same underlying neural mechanisms. More specifically, sequence processing tasks in non-linguistic domains appear to utilise the same area of the brain (the *procedural memory*) as those involved in syntax processing. As navigational planning is one capability needing sequence processing of this type, and was clearly present in pre-linguistic human ancestors, the theory that syntactic language processing evolved as an exaptation of navigational abilities is clearly supported by this evidence.

Two different study techniques were used to demonstrate that linguistic and non-linguistic processing utilised common neural structures. In the first (Hoen et al., 2003),

it was shown that training individuals to resequence data in a non-linguistic task was of clear benefit in helping them to reorder components of a sentence. In the other important experimental study (Hermer-Vasquez et al., 1999), it was shown that while speaking, individuals suffered impairment in another non-linguistic ability, specifically one involved in navigation. The results of both of these experiments are consistent with the idea of common processing devices underlying both linguistic and non-linguistic abilities.

An alternative hypothesis which has not been explored in this thesis is that syntax in language emerged as an exaptation of some other non-linguistic sequence processing task. While the neurological evidence cannot differentiate between this hypothesis and the one pursued in this thesis, the theory of syntax having its roots in navigation and not some other sequence processing task is supported by other arguments. Firstly, on an intuitive level, pre-linguistic human ancestors were certainly capable of navigation; navigational abilities requiring regular language equivalent processing are probably present in even relatively simple creatures such as ants and bees, while rats and other mammals probably have abilities exceeding these capacities. Secondly, Hermer-Vasquez et al. (1999) have shown the presence of a common structure underlying language and orientation; if one aspect of navigation is shown to have a common underpinning with language, it increases the likelihood that another aspect of navigation (i.e. route planning) may also be linked to linguistic processing. Finally, the experimental chapters of this thesis have shown that using syntactic communication for navigation can provide an evolutionary advantage; if syntax-handling in language evolved from pre-existing navigational abilities, it is easy to conceive of the idea that language may have been initially used for this cooperative navigational purposes which promotes its use.

9.1.2 Simulation Design

In chapter 4 it was shown that while existing computer models in this field have much to say about the evolution of languages, most fail to address the evolution of the language faculty. Even in those models in which agents have a genetic component, the evolved

aspect is generally an innate ‘language’, not a language faculty as the term is generally understood in the linguistics community.

As this thesis concerned itself with the evolutionary origins of syntax, it was necessary to construct a new computer model in which the evolution of the language faculty, not of the language, could be studied. In constructing this model, the aim was to assess whether the use of different forms of language could provide benefits that would lead to their promotion and spread by natural selection. In creating a model of this type, it is important to create a system in which language use or competence is not directly rewarded, but may be viewed as a tool for a purpose. If this is not so then the model will almost invariably lead to complex language, producing results which are striking yet inevitable. For this reason, the model developed consisted of a framework which allowed agents with a range of cognitive abilities to undertake the realistic tasks of surviving and reproducing. Any advantages or disadvantages associated with a behaviour can thus be seen in the population size, rather than any measure directly related to the population’s use of language.

An important aspect of the behaviour designed into the model is the formalism developed for navigation. The model was designed to require agents to repeatedly need to visit food and water to survive, and provided a collection of landmarks by which navigation between these resources could be conducted. In such a setting, the most basic form of navigation which permits consistent, reliable navigation between two resources was formally specified and subsequently shown to be regular language equivalent. Importantly, this demonstrates that this basic form of navigation requires that individuals possess the cognitive abilities that would be needed for basic syntax if applied to language. Given that these navigational abilities were presumably present in our ancestors before the evolution of the language faculty occurred, the theory that these abilities were simply exapted by the linguistic system to give syntax, rather than re-evolving from scratch, appears evolutionarily plausible.

While the form of language that the exaptation of such abilities permits demonstrates many characteristics of human language (such as infiniteness, recursion and composi-

tionality), it lacks in complexity compared to modern human language. It is universally agreed that modern human language is at least context-free, yet the theory advanced above, that language is adapted from this simple form of navigation, only leads to a regular language. Two potential solutions to this problem were mentioned. Firstly, early language may have been regular, with context-free abilities only evolving later due to evolutionary pressures which acted on the language faculty. Alternatively, the language faculty may not have exapted the sequence processing capabilities of the regular language equivalent form of navigation, but of a slightly more complex form of navigation which is equivalent to a context-free language. In chapter 5 it was shown that just such a form of navigation would be necessary if one were to need the ability to return along the same path as one took on the outgoing journey. This might be useful for minimising one's exposure to predators, for example.

9.1.3 Experimental Results

In order to establish the primary aim of this thesis, that of showing that syntax in language may have emerged as the result of an exaptation of neural mechanisms used for navigation, it was necessary to show that two theories may be true. The first of these is that this represents a step which is simple enough to be evolutionarily plausible. This was established through a review of the literature and through the formal description of navigational planning as a regular language equivalent task in chapter 5. The second theory is that the existence of syntactic language brings some advantage to its users. In order to prove this, various simulations were run using the model developed in chapter 5.

Observing the impact that different environments had on the size of a population showed that the role of the environment was important. Linguistic populations were seen to form the most viable communities when resources were close together and when these resources were static. However, the same pattern applies for agents who do not communicate. Therefore a measure was required that could establish whether language provided any benefit above the non-communicating baseline behaviour. To this end the *relative benefit* and the *relative improvement* were defined in chapters 6 and 7

respectively. Through the use of these statistics, it was demonstrated that a definite region exists in the studied parameter space for which syntactic language offered a huge benefit over both the selfish behaviour of not communicating and a non-linguistic form of altruism, and in which one would therefore expect strong evolutionary pressure to act for the selection of syntactic language use. This region occurred when either resources were moderately volatile or had a medium degree of spacing. Language use was never seen to be a worse behaviour than the baseline behaviours.

One limitation of this first experiment is that it fails to consider a large range of behaviour types. The use of syntactic language is only compared to the behaviour of agents unable to communicate and thus it is impossible to ascribe the observed benefit to either language use or to the existence of syntax in this language. In order to address this, experiments presented in chapter 7 presented an evaluation of a larger range of navigational and communicative behaviours. These behaviours were chosen such that each added a relatively small additional cognitive ability to the previous, and thus could be evolutionarily plausible. By examining the relative improvements of these behaviours, it was possible to obtain an idea of the evolutionary pathway that would need to be followed to get from silent, wandering behaviour to syntactic communication and route planning. The results revealed a fairly consistent pattern of relative advantage through the parameter space, and showed stronger selective pressure to develop a complex navigational system followed by a communicative system, rather than incrementally adding to both systems simultaneously. This is consistent with the theory presented in this thesis that language may have exapted these sequence processing abilities after they had evolved first for navigation. The results also indicated that a slightly less complex form of communication than fully compositional language was the most advantageous, though this was very likely an artefact of the simulation model.

A question of much interest in the realm of language evolution is why humans are the only creature to have evolved a language system for communication. While the results of these simulations cannot answer this major open question, it does at least suggest a highly speculative new theory. The shift from tree-dwelling apes to hominids which

spent a much greater proportion of their time on Africa's grassland plains is widely acknowledged as one of the major events in the evolution of humans. This shift is associated with such significant changes as our upright posture. While an exact date for the evolution of language is disputed, it is widely considered to have occurred at sometime after this change. During this transition of habitat, the relative importance of different navigation strategies would have changed to require much better navigation over a large range through the use of landmarks, such as in the model presented in this thesis. Given that the genes controlling navigational behaviour were presumably in a state of flux at this time, and other significant changes were being carried out by evolution, it is possible that this unique habitat shift in a social and intelligent species co-occurred with a random mutation which linked the communication system to this regular language equivalent processor. Syntactic language would then emerge as travel over longer distances became more useful and the ability to remember and communicate these directions increased in utility. The experimental results obtained would appear to give tentative backing to this theory; the maximum relative benefit was derived from language use when food was more widely spaced (as would be expected on grasslands), not when it was closer together (as in more heavily forested areas). It is also possible that the benefit associated with more volatile resources may have played a role; it is believed that at this time that our ancestor's diet changed to include a greater proportion of meat, which is a more volatile resource than the plants which had been their primary food source before this.

9.2 Future Directions

The focus of this thesis differs from much of the currently active research in the field of computational language evolution in that it seeks to address the evolution of the language faculty, rather than the evolution of language itself. Broadly speaking, the current trend is for work to seek to explain the features of language in terms of the adaptation that language itself would undergo due to the nature of its acquisition and use. Nevertheless, in these models some assumptions must be made about the

underlying cognitive skills available to the speakers.

While this difference in emphasis makes it difficult to apply the results of this work directly to those of others in this field, it could be hoped that by highlighting the need to explain these basic abilities, awareness of these assumed capabilities in other models will be raised. This should lead others at least to question which low-level abilities they build into their models and whether they can be justified in these situations. Doing so may even lead to these assumptions being stated clearly, so that others can see clearly which axioms the obtained results are predicated upon.

Outside the realm of computational simulation, this thesis also raises questions about the grounding of language and navigation in the brain. In looking for linguistic abilities in non-human creatures, researchers evaluate individuals at their performance on linguistic tasks (see for example Fitch and Hauser (2004)). Having shown in this thesis that some navigational and linguistic skills are computationally equivalent and that they may be based on common neural modules, a new question can be asked in such studies, specifically ‘Do animals have the ability to process information in non-linguistic domains in a way which would be equivalent to syntax if used for communication?’⁴⁵ If the answer to this question for some species is yes, then the issue of why this processing never became available for communication would deserve more attention. If the answer is no, then this suggests that language is uniquely human as other species lacked one component in their ‘readiness for language’.

In looking at possible future directions more directly related to the methods used in this thesis, three aspects of the current model stand out as being open to future refinement. These are the evolution of context-free language and navigation, the coevolution of syntax and lexicon, and the evolution of individual rather than group behaviour. Each of these issues will now be considered in turn.

As stated throughout this thesis, the representation developed for navigation and communication is only equivalent to a regular language, while human language is at

⁴⁵As shown in this thesis, one such domain where these skills may be present and could be tested is in animals’ abilities to carry out regular language and context-free language equivalent navigation.

least context-free. One proposed extension to the representation which would make it context-free would enable agents to navigate to a point and then back along the same route. In the model developed in this thesis, there was no conceivable advantage to doing so and hence the idea was not explored experimentally. Modifications to the system such as the introduction of predators or obstacles in the environment may provide some advantage in returning along the outbound route, and could allow for context-free equivalent navigation and communication to be studied.

The imposition in this model of a lexicon external to the agents is highly unusual in this field, and throughout has been justified by stating that others have shown the coordination of lexicon is possible and hence it can be assumed as pre-existing here. Altering the model to remove this assumption and evolving the lexicon at the same time as developing the syntactic component could lead to some potentially interesting insights. Most trivially, it would be possible to see if the assumption of a pre-existing lexicon used was necessary, or if an evolving lexicon could produce similar results. Further to this, the interaction between the lexical coordination and the benefits of language could be studied. The more coordinated the lexicon, the more beneficial the use of language should be. However, the use of the lexicon to communicate directions means that once two individuals share enough of the lexicon to share directions to a set of resources, they should meet more often and hence coordinate their lexicons further. In this way, one could anticipate that the use that is made of language in this model should create a positive feedback mechanism which will effect both the dynamics of the lexical coordination and of the population through the use of this lexicon.

One final modification to the presented model would be to change the current setting, which corresponds to group selection, into an alternative scheme using the more modern neo-Darwinian framework. This issue was previously addressed in more detail in chapter 5. In the presented work, the measure of the benefit of a behaviour is assessed through its effect on the population. An alternative measure would be to assess how it affected the individual possessing that behaviour; it is not clear from the work presented in this thesis whether it is in an individual's own interest to share knowledge freely, or whether one should remain silent and just observe the speech of others. It is

certain that communication benefits the population (within given environments), but not clear that it benefits the individuals. As shown in chapter 4, the issue of a behaviour which is best for a population being susceptible to freeloaders or liars is well known and studied. The consideration of mechanisms such as kin selection or reciprocal altruism may shed light on whether the results in this thesis are resistant to selfish individuals. This in turn may answer questions as to how important sociality and altruism were as precursors to the evolution of human language.

9.3 Summary

This thesis has argued in favour of a specific account of the genesis of syntax-handling capabilities in humans. This account is that this intrinsic feature of language has evolved as an exaptation of earlier navigation skills. Through surveying current research, it has been argued that this aspect of language may be grounded in the same neural structure as navigational abilities. A formalism was subsequently developed which illustrated that the abilities needed for even quite simple navigation were computationally equivalent to those needed to process regular language. If these abilities were already present in pre-linguistic hominids for navigation, exaptation offers an evolutionarily plausible route by which they may have arisen in language. Finally, computer modelling was used to show that the ability to use syntactic language offers evolutionary benefits to its possessors, thus demonstrating why it may have spread to be available to all modern humans.

References

- Ackley, D. H. and Littman, M. L. (1994). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Artificial Life IV*, pages 40–48, Cambridge, MA. MIT Press.
- Aha, D. W., Kibler, D., and Albert, M. K. (1991). Instance-based learning algorithms. *Machine Learning*, 6(1):37–66.
- Altmann, G. T. M. and Dienes, Z. (1999). Rule learning by seven-month-old infants and neural networks. *Science*, 284:875.
- Arnold, K. and Zuberbuhler, K. (2006). Semantic combinations in primate calls. *Nature*, 441:303.
- Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York.
- Baerends, G. P. (1941). Fortpflanzungsverhalten und orientierung der grabwespe. *Am-mois Jur. Tijdschrift voor Entomologie*, 84:68–275.
- Balaban, E. (1988). Bird song syntax: Learned intraspecific variation is meaningful. *Proceedings of the National Academy of Sciences*, 85:3657–3660.
- Bartlett, M. and Kazakov, D. (2005a). Comparing resource sharing with information exchange in co-operative agents, and the role of environment structure. In Kudenko, D., Kazakov, D., and Alonso, E., editors, *Adaptive Agents and Multi-Agent Systems II, Adaptation and Multi-Agent Learning*, volume 3394 of *Lecture Notes in Artificial Intelligence*, pages 41–54. Springer.
- Bartlett, M. and Kazakov, D. (2005b). The origins of syntax: From navigation to language. *Connection Science*, 17:271–288.
- Bartlett, M. and Kazakov, D. (2006). The evolution of syntactic capacity from navigational ability. In Cangelosi, A., Smith, A. D. M., and Smith, K., editors, *The Evolution of Language*, pages 393–394. World Scientific.
- Barwick, L. and Marrett, A. (2003). Aboriginal traditions. In *Currency Companion to Dance and Music in Australia*, pages 26–28. Currency Press.

-
- Batali, J. (2002). The negotiation and acquisition of recursive grammars as a result of competition among exemplars. In Briscoe, T., editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge University Press.
- Bates, E. and Elman, J. (1996). Learning rediscovered. *Science*, 274:1849–1850.
- Bennett, A. T. D. (1993). Spatial memory in a food storing corvid. *Journal of Comparative Physiology A*, 173:193–207.
- Beugnon, G. and Campan, R. (1989). Spatial memory and spatial cognition in insects. *Etologia*, 1:63–86.
- Bickerton, D. (1975). *Dynamics of a Creole System*. Cambridge University Press, Cambridge.
- Bickerton, D. (1990). *Language and Species*. University of Chicago Press, Chicago Illinois.
- Bickerton, D. (1998). Catastrophic evolution: The case for a single step from protolanguage to full human language. In Hurford, J. R., Studdert-Kennedy, M., and Knight, C., editors, *Approaches to the Evolution of Language*. Cambridge University Press.
- Bickerton, D. (2000). How protolanguage became language. In Knight, C., Studdert-Kennedy, M., and Hurford, J. R., editors, *The Evolutionary Emergence of Language: Social Functions and the Origins of Linguistic Form*. Cambridge University Press.
- Bickerton, D. (2007). Language evolution: A brief guide for linguists. *Lingua*. To Appear.
- Brighton, H. (2003). *Simplicity as a Driving Force in Linguistic Evolution*. PhD thesis, Theoretical and Applied Linguistics, The University of Edinburgh.
- Cangelosi, A. and Harnad, S. (2000). The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution of Communication*, 4:117–142.

REFERENCES

- Cangelosi, A. and Parisi, D. (1998). The emergence of a language in an evolving population of neural networks. *Connection Science*, 10(2):83–97.
- Cangelosi, A. and Parisi, D., editors (2001). *Simulating the Evolution of Language*. Springer Verlag, London, UK.
- Cartwright, B. A. and Collett, T. S. (1983). Landmark learning in bees. *Journal of Comparative Physiology*, 151:521–543.
- Chapula, N. and Scardigli, P. (1993). Shortcut ability in hamsters (*Mesocricetus auratus*): The role of environmental and kinesthetic information. *Animal Learning and Behavior*, 21:255–265.
- Cheney, D. L. and Seyfarth, R. M. (1990). *How Monkeys See the World: Inside the Mind of Another Species*. University of Chicago Press, Chicago.
- Cheng, K. (1986). A purely geometric module in the rat’s spatial representation. *Cognition*, 23:149–178.
- Cheng, K., Collett, T. S., Pickhard, A., and Wehner, R. (1987). The use of visual landmarks by honeybees: Bees weight landmarks according to their distance from the goal. *Journal of Comparative Physiology A*, 161:469–475.
- Cheng, K. and Sherry, D. (1992). Landmark-based spatial memory in birds (*Parus atricapillus* and *Columba livia*): The use of edges and distances to represent spatial positions. *Journal of Comparative Psychology*, 106:331–341.
- Chittka, L. and Geiger, K. (1995). Honey-bee long-distance orientation in a controlled environment. *Ethology*, 99:117–126.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. MIT Press, Cambridge, Massachusetts.
- Chomsky, N. (1980). *Rules and Representations*. Basil Blackwell, London.
- Chomsky, N. (1986). *Knowledge of Language*. New York: Praeger.

-
- Chomsky, N. (1988). *Language and Problems of Knowledge*. MIT Press, Cambridge, Massachusetts.
- Chomsky, N. (1995). *The Minimalist Program*. MIT Press, Cambridge, Massachusetts.
- Clark, R., Gleitman, L., and Kroch, A. (1997). Acquiring language: Letter to Science. *Science*, 276:1177–1181.
- Collett, T., Cartwright, B. A., and Smith, B. A. (1986). Landmark learning and visuo-spatial memories in gerbils. *Journal of Comparative Physiology A*, 6:835–851.
- Collett, T. S. (1996). Insect navigation en route to the goal: Multiple strategies for the use of landmarks. *Journal of Experimental Biology*, 199:227–235.
- Collett, T. S., Dittman, E., Giger, A., and Wehner, R. (1992). Visual landmarks and route-following in desert ants. *Journal of Comparative Physiology A*, 170:435–442.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, Oxford.
- Denebourg, J. L., Pasteels, J. M., and Vanhaeghe, J. C. (1983). Probabilistic behaviour in ants: A strategy of errors? *Journal of Theoretical Biology*, 105:259–271.
- Dienes, Z., Altmann, G. T. M., and Gao, S.-J. (1999). Mapping across domains without feedback: A neural network model of transfer of implicit knowledge. *Cognitive Science*, 23(1):53–82.
- Dorigo, M., Caro, G. D., and Gambardella, L. M. (1999). Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172.
- Dyer, F. C. (1991). Bess acquire route-based memories but not cognitive maps in a familiar landscape. *Animal Behaviour*, 41:239–246.
- Dyer, F. C. (1994). Spatial cognition and navigation in insects. In Real, L., editor, *Behavioural Mechanisms in Evolutionary Ecology*, pages 66–98. Chicago University Press, Chicago.
- Dyer, F. C. and Dickinson, J. A. (1996). Sun-compass learning in insects: Representation in a simple mind. *Current Directions in Psychological Science*, 5:67–71.

REFERENCES

- Eichenbaum, H. (2000). A cortical-hippocampal system for declarative memory. *Nature Review Neuroscience*, 1(1):41–50.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14:179–211.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48:71–99.
- Etienne, A. S., Maurer, R., and Saucy, F. (1988). Limitations in the assessment of path dependent information. *Behavior*, 106:81–111.
- Fitch, W. T. and Hauser, M. D. (2004). Computational constraints on syntactic processing in a nonhuman primate. *Science*, 303:377–380.
- Foster, D. J. and Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440:680–683.
- Gallistel, C. R. (1990). *The Organisation of Learning*. MIT Press, Cambridge, Massachusetts.
- Gallistel, C. R. (1994). Space and time. In Mackintosh, N. J., editor, *Animal Learning and Cognition*, pages 221–253. Academic Press, San Diego.
- Geissmann, T. (2002). Duet-splitting and the evolution of gibbon song. *Biological Reviews*, 77:57–76.
- Gentner, T. Q., Fenn, K. M., Margoliash, D., and Nusbaum, H. C. (2006). Recursive syntactic pattern learning by songbirds. *Nature*, 440:1204–1207.
- Goodale, M. A., Ellard, C. G., and Booth, L. (1990). The role of image size and retinal motion in the computation of absolute distance by the mongolian gerbil (*Meriones unguiculatus*). *Vision Research*, 30:399–413.
- Goschke, T., Friederici, A. D., Kotz, S. A., and van Kampen, A. (2001). Procedural learning in Broca’s aphasia: Dissociation between the implicit acquisition of spatio-motor and phoneme sequences. *Journal of Cognitive Neuroscience*, 13(3):370–388.
- Gould, J. L. (1982). The magnetic sense of pigeons. *Nature*, 296:205–211.

- Gould, J. L. (1986). The locale map of honey bees: Do bees have a cognitive map? *Science*, 232:861–863.
- Gould, S. J. and Lewontin, R. C. (1979). The spandrels of San Marco and the panglossian program: A critique of the adaptationist programme. *Proceedings of the Royal Society of London*, 205:281–288.
- Gould, S. J. and Vrba, E. S. (1982). Exaptation – a missing term in the science of form. *Paleobiology*, 8:4–15.
- Graham, P., Durier, V., and Collett, T. S. (2004). The binding and recall of snapshot memories in wood ants (*Formica rufa* L.). *Journal of Experimental Biology*, 207:393–398.
- Graham, P., Fauria, K., and Collett, T. S. (2003). The influence of beacon-aiming on the routes of wood ants. *Journal of Experimental Biology*, 206:535–541.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour I and II. *Journal of Theoretical Biology*, 7:1–32.
- Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics SSC4*, 2:100–107.
- Hashimoto, T. and Ikegami, T. (1996). Emergence of net-grammar in communicating agents. *BioSystems*, 38:1–14.
- Hashimoto, T. and Nakatsuka, M. (2006). Reconsidering Kirby’s compositionality model toward modelling grammaticalisation. In Cangelosi, A., Smith, A. D. M., and Smith, K., editors, *The Evolution of Language*. World Scientific.
- Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). The faculty of language: What is it, who has it and how did it evolve? *Science*, 298:1569–1579.
- Hermer, L. and Spelke, E. S. (1996). Modularity and development: The case of spatial reorientation. *Cognition*, 61:195–232.

REFERENCES

- Hermer-Vasquez, L., Spelke, E. S., and Katsnelson, A. S. (1999). Sources of flexibility in human cognition: Dual-task studies of space and language. *Cognitive Psychology*, 39:3–36.
- Hermer-Vazquez, L. (1997). *Cognitive Flexibility as it Emerges over Evolution and Human Development: The Case of Spatial Reorientation*. PhD thesis, Cornell University.
- Hoen, M., Golembiowski, M., Guyot, E., Deprez, V., Caplan, D., and Dominey, P. F. (2003). Training with cognitive sequences improves syntactic comprehension in agrammatic aphasics. *NeuroReport*, 14(3):495–499.
- Hurford, J. (2000). Social transmission favours linguistic generalization. In Chris Knight, J. R. H. and Studdert-Kennedy, M., editors, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pages 324–352. Cambridge University Press, Cambridge.
- Hurford, J. (2002). Expression/induction models of language evolution: Dimensions and issues. In Briscoe, T., editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, chapter 10. Cambridge University Press.
- Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press.
- Jacob, F. (1977). Evolution and tinkering. *Science*, 196:1161–1166.
- Janik, V. M. (2000). Whistle matching in wild bottlenose dolphins (*Tursiops truncatus*). *Science*, 289:1355–1357.
- Kazakov, D. and Bartlett, M. (2002). A multi-agent simulation of the evolution of language. In Grobelnik, M., Bohanec, M., Mladenic, D., and Gams, M., editors, *Proceedings A of the Fifth International Multi-Conference Information Society (IS 2002)*, pages 39–41, Ljubljana, Slovenia. Josef Stefan Institute.
- Kazakov, D. and Bartlett, M. (2004a). Co-operative navigation and the faculty of language. *Applied Artificial Intelligence*, 18(9–10):885–901.

- Kazakov, D. and Bartlett, M. (2004b). Social learning through evolution of language. In Liardet, P., Collet, P., Fonlupt, C., Lutton, E., and Schoenauer, M., editors, *Artificial Evolution, 6th International Conference, Evolution Artificielle, EA 2003*, volume 2936 of *Lecture Notes in Computer Science*, pages 397–408. Springer.
- Kazakov, D. and Bartlett, M. (2005). Could navigation be the key to language? In *Proceedings of the Second Symposium on the Emergence and Evolution of Linguistic Communication (EELC'05)*, pages 50–55, Hatfield, UK.
- Kimura, D. (1999). *Sex and Cognition*. MIT Press, Cambridge, Massachusetts.
- Kingsolver, J. C. and Koehl, M. A. R. (1985). Aerodynamics, thermoregulation, and evolution of insect wings: Differential scaling and evolutionary change. *Evolution*, 39:448–504.
- Kirby, S. (1998). Fitness and the selective adaptation of language. In Hurford, J. R., Studdert-Kennedy, M., and Knight, C., editors, *Approaches to the Evolution of Language: Social and Cognitive Bases*. Cambridge University Press.
- Kirby, S. (2000). Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. In Knight, C., editor, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pages 303–323. Cambridge University Press.
- Kirby, S. (2002a). Learning, bottlenecks and the evolution of recursive syntax. In Briscoe, T., editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge University Press.
- Kirby, S. (2002b). Natural language from artificial life. *Artificial Life*, 8(2):185–215.
- Krebs, J. R. and Dawkins, R. (1984). Animal signals: Mind-reading and manipulation. In Krebs, J. R. and Davies, N. B., editors, *Behavioural ecology. An evolutionary approach*, pages 380–402. Blackwell.
- Loyacano, H. A. and Chappell, J. A. (1977). Sun-compass orientation in juvenile large-mouth bass *Micropterus salmoides*. *Transactions of the American Fisheries Society*, 106:77–79.

REFERENCES

- MacLennan, B. and Burghardt, G. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–188.
- Maki, P. M. and Resnick, S. M. (2000). Longitudinal effects of estrogen replacement therapy on pet cerebral blood flow and cognition. *Neurobiology of Aging*, 21(2):373–383.
- Marcus, G. F., Vijayan, S., Rao, S. B., and Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, 283:77–80.
- Maurer, R. and Seguinot, V. (1995). What is modelling for? A critical review of the models of path integration. *Journal of Theoretical Biology*, 175:457–475.
- McClelland, J. L. and Plaut, D. C. (1999). Does generalization in infant learning implicate abstract algebra-like rules? *Trends in Cognitive Sciences*, 3:166–168.
- Miles, C., Green, R., Sanders, G., and Hines, M. (1998). Estrogen and memory in a transsexual population. *Hormones and Behavior*, 34(2):199–208.
- Mishkin, M., Malamut, B., and Bachevalier, J. (1984). Memories and habits: Two neural systems. In Lynch, G., McGaugh, J. L., and Weinburger, N. W., editors, *Neurobiology of Learning and Memory*, pages 65–77. Guilford Press, New York.
- Mittelstaedt, H. and Mittelstaedt, M. L. (1982). Homing by path integration. In Papi, F. and Wallraff, H. G., editors, *Avian Navigation*, pages 290–297. Springer-Verlag.
- Morris, R. G. M. (1981). Spatial localisation does not require the presence of local cues. *Learning and Motivation*, 12:239–260.
- Nehaniv, C. L. (2005). Open problems in the emergence and evolution of linguistic communication: A road-map for research. In Cangelosi, A., editor, *Proceedings of the Second International Symposium on the Emergence and Evolution of Linguistic Communication (EELC'05)*, pages 86–93.
- Newbold, P. (1996). *Statistics for Business and Economics*. Prentice Hall, London, UK, 4th edition.

- Newport, E. H., Gleitman, H., and Gleitman, L. R. (1977). Mother, I'd rather do it myself. In Snow, C. E. and Ferguson, C. A., editors, *Talking to Children: Language Input and Acquisition*. Cambridge University Press.
- Nicholson, D. J., Judd, S. P. D., Cartwright, B. A., and Collett, T. S. (1999). Learning walks and landmark guidance in wood ants (*Formica rufa*). *The Journal of Experimental Biology*, 202:1831–1838.
- Nissen, M. J. and Bullemer, P. T. (1987). Attentional requirements for learning: Evidence from performance measures. *Cognitive Psychology*, 19:1–32.
- Noble, J. (1999). Cooperation, conflict and the evolution of communication. *Adaptive Behavior*, 7(3/4):349–370.
- O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford.
- Oliphant, M. and Batali, J. (1997). Learning and the emergence of coordinated communication. *The Newsletter of the Center for Research in Language*, 11(1).
- Paul, U. S. (1982). Do geese use path integration for walking home? In Papi, F. and Wallraff, H. G., editors, *Avian Navigation*, pages 298–307. Springer-Verlag.
- Pearce, J. M. (1997). *Animal Learning and Cognition: An Introduction*. Psychology Press, Hove, UK, 2nd edition.
- Perruchet, P. and Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences*, 10:233–238.
- Piattelli-Palmarini, M. (1989). Evolution, selection and cognition: From learning to parameter setting in biology and in the study of language. *Cognition*, 31:1–44.
- Pinker, S. (1997). Acquiring language: Letter to Science. *Science*, 276:1177–1181.
- Pinker, S. and Bloom, P. (1990). Natural language and natural selection. *Behavioural and Brain Sciences*, 13:707–784.

REFERENCES

- Pullum, G. K. and Scholz, B. C. (2002). Empirical assessment of stimulus poverty arguments. *The Linguistic Review*, 19:9–50.
- Reese, E. S. (1989). Orientation behaviour of butterflyfishes (family *Chaetodontidae*) on coral reefs: Spatial learning of route specific landmarks and cognitive maps. *Environmental Biology of Fishes*, 25:79–86.
- Saffran, J. R., Aslin, R. N., and Newport, E. L. (1996). Statistical learning by 8-month old infants. *Science*, 274:1926–1928.
- Sampson, G. (1997). *Educating Eve: ‘The Language Instinct’ debate*. Cassell, London.
- Senghas, A. and Coppola, M. (2001). Children creating language: How Nicaraguan Sign Language acquired a spatial grammar. *Psychological Science*, 12:323–328.
- Sherry, D. F. and Duff, S. J. (1996). Behavioural and neural bases of orientation in food-storing birds. *Journal of Experimental Biology*, 199:165–171.
- Sherwin, B. B. (1998). Estrogen and cognitive functioning in women. *Proceedings of the Society for Experimental Biology and Medicine*, 217(1):17–22.
- Shettleworth, S. (1998). *Cognition, Evolution and Behavior*. Oxford University Press, New York.
- Smith, A. D. M. (2003a). *Evolving Communication through the Inference of Meaning*. PhD thesis, Theoretical and Applied Linguistics, School of Philosophy, Psychology and Language Sciences, The University of Edinburgh.
- Smith, A. D. M. (2005). Stable communication through dynamic language. In Cangelosi, A., editor, *Proceedings of the Second International Symposium on the Emergence and Evolution of Linguistic Communication (EELC’05)*, pages 135–142.
- Smith, K. (2003b). *The Transmission of Language: Models of Biological and Cultural Evolution*. PhD thesis, Theoretical and Applied Linguistics, School of Philosophy, Psychology and Language Sciences, The University of Edinburgh.
- Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life*, 2(3):319–332.

-
- Steels, L. (1996a). Perceptually grounded meaning creation. In Tokoro, M., editor, *Proceedings of the International Conference on Multi-Agent Systems*. MIT Press.
- Steels, L. (1996b). Self-organizing vocabularies. In Langton, C. and Shimohara, T., editors, *Artificial Life V*, pages 179–184, Nara, Japan.
- Steels, L. (1997). Language learning and language contact. In Daelemans, W., Van den Bosch, A., and Weijters, A., editors, *Proceedings of the Workshop on Empirical Approaches to Language Acquisition*, pages 11–24, Prague.
- Steels, L. and Kaplan, F. (1999). Situated grounded word semantics. In Dean, T., editor, *International Joint Conference on Artificial Intelligence (IJCAI99)*. Morgan Kaufmann Publishers.
- Steels, L. and Vogt, P. (1997). Grounding adaptive language games in robotic agents. In Harvey, I. and Husbands, P., editors, *European Conference on Artificial Life (ECAL97)*, Cambridge, MA. MIT Press.
- Tinbergen, N. (1951). *The Study of Instinct*. Oxford University Press, Oxford.
- Tolman, E. C., Ritchie, B. F., and Kalish, D. (1946). Studies in spatial learning. I. Orientation and the short-cut. *Journal of Experimental Psychology*, 1:13–24.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57.
- Turner, H. and Kazakov, D. (2003). Stochastic simulation of inherited kinship-driven altruism. In Alonso, E., Kudenko, D., and Kazakov, D., editors, *Adaptive Agents and Multi-Agent Systems*. Springer.
- Ugolini, A. (1987). Visual information acquired during displacement and initial orientation in *Polistes gallius* hymenopteron vespidae. *Animal Behaviour*, 35(2):590–595.
- Ullman, M. T. (2004). Contributions of neural memory circuits to language: The declarative/procedural model. *Cognition*, 92:231–270.

REFERENCES

- Vallortigara, G., Zanforlin, M., and Pasti, G. (1990). Geometric modules in animals' spatial representations: A test with chicks *Gallus gallus domesticus*. *Journal of Comparative Psychology*, 104:248–254.
- von Frisch, K. (1967). *The Dance Language and Orientation of Bees*. Belknap Press of Oxford University Press, Cambridge, Massachusetts.
- Wagner, K., Reggia, J., Uriagereka, J., and Wilkinson, G. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69.
- Walcott, C. (1978). Anomalies in the Earth's magnetic field increase the scatter of pigeons vanishing bearings. In Schmidt-Koenig, K. and Keeton, W. T., editors, *Animal Migration, Navigation and Homing*, pages 143–151. Springer.
- Wehner, R. (1992). Arthropods. In Papi, F., editor, *Animal Homing*, pages 45–144. Chapman and Hall.
- Wehner, R. and Lanfranconi, B. (1981). What do the ants know about the rotation of the sky? *Nature*, 293:731–733.
- Werner, G. and Dyer, M. (1992). Evolution of communication in artificial organisms. In Langton, C., Taylor, C., Farmer, D., and Rasmussen, S., editors, *Artificial Life II*, pages 659–687, Redwood City, CA. Addison-Wesley Pub.
- Wiltschko, W. and Balda, R. P. (1989). Sun compass orientation in seed-caching scrub jays (*Aphelocoma coerulesceus*). *Journal of Comparative Physiology A*, 164:717–721.
- Wiltschko, W. and Wiltschko, R. (1991). Magnetic orientation and celestial cues in migratory orientations. In Berthold, P., editor, *Orientation in Birds*, pages 16–37. Birkhauser Verlag.
- Wittlinger, M., Wehner, R., and Wolf, H. (2006). The ant odometer: Stepping on stilts and stumps. *Science*, 312:1965–1967.
- Wray, A. (2005). Needs only analysis in linguistic ontogeny and phylogeny. In Cangelosi,

- A., editor, *Proceedings of the Second International Symposium on the Emergence and Evolution of Linguistic Communication (EELC'05)*.
- Zahavi, A. (1975). Mate selection - a selection for a handicap. *Journal of Theoretical Biology*, 53:205–213.
- Zeil, J., Kelber, A., and Voss, R. (1996). Structure and function of learning flights in bees and wasps. *Journal of Experimental Biology*, 199:245–252.
- Zuidema, W. and Hogeweg, P. (2000). Selective advantages of syntactic language - a model study. In *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, pages 577–582, Hillsdale, USA. Lawrence Erlbaum Associates.