

LEARNING AND THE EVOLUTION OF LANGUAGE: THE ROLE OF CULTURAL VARIATION AND LEARNING COST IN THE BALDWIN EFFECT

Steve Munroe

Intelligent Agent Group
Department of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ
sjm01r@ecs.soton.ac.uk

Angelo Cangelosi

Centre for Neural and Adaptive Systems
School of Computing
University of Plymouth
Plymouth PL4 8AA
acangelosi@plymouth.ac.uk

Abstract

The Baldwin Effect has been explicitly used as an argument for the explanation of the origins of language and the language acquisition device (Pinker & Bloom, 1990). This paper presents new simulations of an artificial life model for the evolution of compositional languages (Cangelosi, 2001). It specifically addresses the role of cultural variation and of learning costs in the Baldwin effect for the evolution of language. Results shown that when there is a high cost associate to language leaning, agents gradually assimilate in their genome some explicit features (e.g lexicon) of the specific language they are exposed to. When the structure of the language is allowed to vary during cultural transmission, Baldwinian processes cause the assimilation of a predisposition to learn, rather than assimilating any structural properties associated with a specific language. The analysis of the mechanisms underlying such language learning/evolution effects (e.g. categorical perception) supports Deacon's (1997) hypothesis on the Baldwinian inheritance of general underlying cognitive capabilities, rather than assimilation of a language acquisition device's principles and parameters.

Keywords: Baldwin effect, language evolution, neural networks, genetic algorithms, learning cost, cultural transmission

Short title: Baldwin effect in language evolution

LEARNING AND THE EVOLUTION OF LANGUAGE: THE ROLE OF CULTURAL VARIATION AND LEARNING COST IN THE BALDWIN EFFECT

1. Introduction

The interaction between evolution and learning, and phenomena such as the Baldwin effect, are of fundamental importance in artificial life systems that use both evolutionary computation and learning. A series of computational models have directly investigated the way evolution affects learning, and the effects of learning on the evolution of adaptive traits. Hinton & Nowlan's (1987) model of the Baldwin effect is one of the most influential papers that investigated the process of the assimilation of learned traits in genotypes. Other computer simulations studies (Turney et al., 1996; Belew & Mitchell, 1986) have investigated in detail some of the Baldwin effect mechanisms. In addition, other research (Nolfi and Floreano, 1999; Mayley, 1996) has studied the various mechanisms of interaction between learning and evolution and the effects of different costs attached to learning.

The Baldwin effect is a 100-year-old argument about learning and evolution. It studies the fact that what a species must initially learn during each member's lifetime can over time become part of the genetic makeup of that species. In other words, what initially needed to be learnt eventually becomes innate. This argument was first proposed by the 19th Century theorist James Mark Baldwin (Baldwin, 1896). Superficially it sounds like a variant of Lamarckism, the discredited theory of the evolution of acquired characteristics. In reality, Baldwinian evolution sits uncontroversially within the traditional Neo-Darwinian evolutionary framework. Baldwin's argument runs as follows: Imagine a species that must learn a certain task if it is to be successful in its environment. Now, a typical example of that species must initially start from scratch in learning this task, i.e. it has no innate ability. Also learning has costs. It could be dangerous to perform the task poorly and the time and energy it takes to learn the task impacts on other important activities such as looking for food etc. Furthermore, while attempting to learn the task the individual could be distracted from noticing prey or predators and so on. Over generations some individuals will be born such that through chance mutations in their

genome they have a small amount of innate ability to perform the task. Perhaps their reactions are a little quicker, or the neural mechanisms that are used to perform the task are a little more efficient. These individuals will learn the task more quickly and will also perform it better. This will increase their fitness and hence their reproductive ability. Thus over time this process works like a ratchet effect, slowly pushing the species closer to specific behavior.

The study of the relationship between acquired traits and innate constraints, and in general the nurture/nature debate, is of great importance also in cognitive science (Elman et al., 1996). This also applies in the field of the acquisition and evolution of language. For example, the Baldwin Effect has been explicitly used as an argument for the explanation of the origins of language (e.g. Pinker & Bloom, 1990). Central to the study of language is the question of how much does language development owe to innate characteristics. Debate in this area has continued endlessly. At the one extreme the behaviorists argued that children learn language through the same associationist learning mechanism with which they learn everything else, and as such language presents us with no special problem (Skinner, 1953). Though this position is no longer widely held, there is no question that children do learn much of language through observation and explicit teaching from parents, as the behaviorists believe. Indeed, some recent theories of language acquisition stress the essential role of learning, such as in constructivist theories of syntax acquisition (Tomasello & Brooks, 1999).

At the other end of the spectrum there are those that point out that the ease with which children learn language suggests some specialized language learning mechanism in the Brain. Chomsky's idea of the Language Acquisition Device (LAD) was put forward as just such a mechanism. Chomsky argued that when growing up, the examples children get of language use from adults could never amount to a coherent and systematic enough data set from which children could become skilled language users (Chomsky, 1975). This poverty of stimulus argument was put forward as evidence as to why it could be reasonable to assume some specialized device in the brain - the LAD - that would do all the hard work in learning a language. But how much work can be attributed to the LAD and how much to the learning abilities of the child? This question is of major importance in linguistics and the study of language in general.

Traditionally, linguists have been pessimistic towards arguments favouring an evolutionary account of language. This stems from Chomsky's argument that the form a language takes has

little bearing on the fitness of the language-using individual. Thus evolutionary processes could not have sculpted the shape or form of language. For example it seems difficult to accept that the particular way a language expresses plurality could have any adaptive value; the ability to express plurality could be seen to be adaptive but surely the details of its implementation could not so easily be seen to be adaptive. However, more recently theorists have begun to look to evolution and the Baldwin effect as possible ways to account for language evolution, and the old Chomskyan pessimism no longer holds sway like it did in previous years. However there is still argument over just what could have been evolved through a Baldwinian process with regards to language. Some argue for a fully specified LAD, which sets out the 'principles and parameters' for all languages whilst others argue that only some substrates of cognition that underlie language ability (but nevertheless are not language-like) could have evolved through Baldwinian evolution (e.g. Deacon, 1997).

Recently, computational models have been employed to simulate the evolution of language (Cangelosi & Parisi, 2002; Kirby, in press). Some models simulate both language learning (i.e. cultural transmission) and evolution (i.e. inheritance of linguistic traits), therefore these can be useful for investigating how the two mechanisms interact. In this paper, we will present new simulations on a model of the evolution of compositional languages that specifically address the role of the cultural transmission and of the learning cost in the Baldwin effect. In the next sections, we will first look at the contribution of such models in understanding the Baldwin effect and in general the interaction of learning and evolution. Subsequently, we analyze the current modeling research on the Baldwin effect and the evolution of language. Learning is here represented by the process of the cultural transmission of language (glossogenesis). We then present new simulations that further explore the relation between the evolution of language, the role of cultural transmission, and the effects of varying learning costs in the Baldwin effect.

2. Modeling the Interaction of Learning and Evolution

The most influential study of the Baldwin effect is that of Hinton and Nowlan (1987). They used a Genetic Algorithm (GA) to display how the Baldwin effect takes place. In this model, the individuals that make up a population have a genome represented as a string of 0, 1 and a special wildcard allele ?, which represents a position on the genome that could be either a 0 or

a 1 (depending on a simple form of learning, i.e. a random guess). Hinton and Nowlan constructed a fitness landscape that had only one point surrounded by a completely flat space. In this situation the individuals all have zero fitness unless they are on top of the point. They used two different kinds of individuals in the experiment: those that can learn during their lifetime and those that cannot. It was shown that the individuals that can learn effectively smoothed out the landscape enabling the GA to improve performance. This ability translated into the individual's genome and as such is a good example of Baldwinian Mechanisms at work.

In a more recent paper, Stefano Nolfi and Dario Floreano (1999) argue that by merging the processes of intra generational learning (lifetime adaptation) and inter generational learning (evolutionary adaptation) it is possible to develop systems that can adapt better to dynamic environments than systems designed with only one such process. They discuss a series of examples of the adaptive function of learning in evolution. For example, learning allows organisms to adapt to “quick” changes, i.e. those that happen within the organism’s lifetime, or in a few generations (e.g. simulations of Miller & Todd, 1991). The Baldwin effect, i.e. when learning “guides” evolution for the assimilation of learned traits, is one such adaptive function.

Among the many models analyzed in Nolfi & Floreano (1999), one focuses on the effects learning can have on evolution if the fitness spaces of the evolutionary task (i.e. forage) and of the learning task (i.e. predict next sensory input) are dynamically correlated. This dynamic correlation refers to the fact that when there is (local) improvement in the learning task, this results in some improvement also for the evolutionary task. If the landscapes of both learning and evolutionary spaces are dynamically correlated, then a plastic individual can through learning, bring itself into a higher position on the evolutionary landscape (Nolfi, Elman & Parisi, 1994). This is also true when the two tasks are clearly uncorrelated (Harvey, 1997; Nolfi, 2000). Another model shows that sometimes Baldwinian processes operate at a more subtle level than simply the direct assimilation of aspects of behavior. In Nolfi and Parisi’s (1997) model, simple khepera robots have to find a target in an arena whose walls change color. In parallel they perform a self-teaching task. When comparing populations of robots that learn during their lifetime and those that do not learn, results showed that the first perform better than the latter, and that learning robots acquire characters that are particularly adapted to the environment they spend most time in. More importantly, it was shown that the robots belonging to the learning populations have not simply assimilated particular environment-specific behavior, rather they have coded a *predisposition to learn* to behave efficiently. In other

words, Baldwinian processes have used information gained through the succeeding generations' ability to engage in lifetime learning, to assimilate an innate capacity that will facilitate learning. These results are very relevant for the study of Baldwinian effects in the evolution of language. Different useful traits for language can be acquired, some for assimilating the specific lexicon at birth, other for assimilating other structures useful for the specific language to be learned after birth.

Some important issues that affect the interaction between learning and evolution are the actual feasibility of phenomena such as the Baldwin effect, and the trade-off between benefits and costs of learning. In fact, some environmental features can sometimes be hard, if not impossible, to learn and/or assimilate. Ackley & Littman (1991) state that there is an upper limit in that too much variation makes them unpredictable and therefore unlearnable. In other cases, learning may help to generate behaviors that evolution would find hard to evolve. The precise structure of the English language for example is not likely to have been specified by evolution, but rather some general principles and parameters may have been genetically specified whereas the details of the language are likely to be left for the user to learn (Pinker, 1994). One of the main aims of this project is to see exactly what can be specified in terms of language by Baldwinian processes.

The trade-off between benefits and costs of learning also affects the feasibility of Baldwinian phenomena. Whilst it seems clear that there are some benefits for a species to display some plasticity in its interaction with the environment, it is perhaps less clear that there will be some costs also. Mayley (1996) outlines four of these costs. First there is a cost in terms of the length of time necessary to learn the specific behavior. If the behavior were innate, then there would not be delay in fitness optimization. Also, in species with this kind of learning there is usually a need for parental care in the early years of life. Secondly, whilst learning the behaviour there will be many instances of incorrect or poor performance. Depending on the behaviour this could be dangerous and as such could impact on the fitness of the individual. The third cost is in the genetic complexity required to implement a system of regulatory processes that guide the learning mechanisms as opposed to a simple fixed structure necessary for innate behaviour. Finally, there is a cost in the learning itself in that the time spent learning must be subtracted from other important activities such as looking for food, mating, or predator avoidance. Given these costs Mayley argues that learning will only be used when necessary,

e.g. when an agent is in such an environment that the ability to learn is essential for the agent's fitness, for example when the environment is changing more rapidly than evolution can track.

Mayley investigated the costs of learning in various experiments. He specified two kinds of costs. The first - explicit costs - refers to situations in which the cost is discontinuous with the learning itself; this is analogous to supervised learning in artificial neural networks, where the backpropagation of error is not computed until a run on the learning task has been completed. In the natural world this kind of costs refers to situations in which the individual performs the behavior wrongly and then pays some price for the mistake, i.e. injury or death. The study that looked at this kind of cost found that it provided enough motivation for the system to develop genetic assimilation of learnt traits. The second kind of cost is termed implicit and refers to cases where the cost of learning is contiguous to the actual process of learning. This is analogous to the time and energy consumption costs mentioned above and is also similar to unsupervised learning in neural networks. This kind of cost likewise helps to drive the Baldwin effect but has a much stronger effect. Mayley also conducted an experiment where there was no cost associated with learning. In this scenario there was no benefit to genetic assimilation and as such it did not occur. Therefore it can be seen that for Baldwinian processes to occur there must be a downside to learning, i.e. costs of some kind. When these costs are outweighed by the benefits, the evolution of learning will occur, when this situation is inverted then the ability to learn will be gradually assimilate in the genotype.

3. The Baldwin Effect in Language Evolution: Theories and Models

In the influential 1990 *Behavioral and Brain Sciences* paper, Pinker & Bloom put language back into the arena of evolution. This was necessary due to the influence of Chomsky's pessimism in looking at an evolutionary explanation of language. Chomsky argued against a functional account of language, claiming instead that its fundamental property existed in its forms not in its functionality. Furthermore, the form a language takes can have no relation to the fitness of its speakers. For example the characteristic form of any particular language seems to confer no evolutionary advantage to its population of speakers. Populations of language speakers do not compete for survival in terms of the form of their language. Having arrived at this position it seemed improbable to suppose that evolution could have had a say in the design of the form of language.

Pinker and Bloom argue against the two theories used to support Chomsky's ideas: "exaptation" and the "spandrel" theory. Exaptation refers to the situation when a structure or behavior evolved for a particular use is used for another purpose to which by sheer accident it is well designed. For example ears have the useful property of propping up spectacles though they cannot seriously be thought to have been evolved for that purpose, rather their design has been exapted for this use. Similarly language can be argued to have developed from other evolved structures and behaviors without itself ever having an evolutionary trajectory. Spandrel theory refers to the idea that certain structures or behaviors are necessary for the support of other structures or behaviors but in themselves have no direct evolutionary 'purpose'. Here language can be seen as exploiting structural or behavioural aspects of human evolutionary 'design' but remaining invisible to any direct evolutionary processes. Instead Pinker and Bloom argue that as language displays characteristics of a complex design for some function it seems unreasonable to suppose that language was not designed by the only candidate known to science that can create a complex design in order to meet some function, viz. natural selection (cf. also Pinker, 1994).

Using a more neurally-based approach, Terrence Deacon (1997) offers a more biologically plausible account of language origins than what had gone before. Deacon looks at the possibility that language evolved via Baldwinian processes. In opposition to Pinker & Bloom, Deacon persuasively argues against the possibility of the evolution of a Universal Grammar (UG) in that for it to have been translated into the genome, important aspects of it would have to have remained constant for hundreds of generations. Furthermore, for any genetic assimilation to occur the neural substrates that processes language function must remain invariant regardless of the language or the person, again through large numbers of generations. However, the neural aspects of speech production and comprehension have been found to vary widely from person to person. No two people have exactly the same pattern of neural activation across language tasks. In fact the same person can show different activation patterns when performing the same task at different times (Gazzaniga, Ivry & Mangun, 1998). The different surface implementation of UG grammar in different languages makes a genetic account of UG unlikely. Deacon suggests that those aspects of language many linguists would rank most likely to be part of a Universal Grammar are precisely those that are ineligible to participate in Baldwinian evolution. Deacon states categorically the kinds of language properties that cannot participate: "*No innate rules, no general principles, no innate symbolic categories can be built*

by evolution” (Deacon 1997:339). However, Deacon goes on to propose that there are various cognitive pre-dispositions supporting language that can result from Baldwinian processes, for example such cognitive abilities as attention, imitation or automatic reflection. These abilities taken as a whole prepare a groundbed for language acquisition. Thus Baldwinian evolution is seen not to work on language directly but on the processes that underlie a language capability. Deacon goes on to argue that the most likely candidate for the location of the supporting abilities that underpin language is the Pre-Frontal Cortex.

Theories and hypothesis on the evolution of language, such as those on the role of the Baldwin effect, are difficult to test due to the limited availability of data and direct experimental evidence. However, the recent development of computational models for the evolution of language and communication has significantly contributed to the renewed interest in the use of a scientific approach in this field (Cangelosi & Parisi, 2002). These simulation-based studies use a variety of computational methods, mostly based on artificial life and synthetic approaches (Kirby, in press; Parisi, 1997; Steels, 1997). They are used to test specific hypotheses on language origins, or to explore plausible scenarios that lead to the evolution of language and communication. Indeed, current simulation studies of language evolution have dealt with a variety of research issues including the emergence of syntax and syntactic universals (e.g. Brighton, 2002; Cangelosi, 2001; Hazlehurst and Hutchins, 1998; Kirby, 2001; Nowak, Plotkin & Jansen, 2000; Teal & Taylor, 2000), simulations that focus on the evolution and auto-organization of communication systems (Arita & Taylor, 1996; Cangelosi & Parisi, 1998; de Boer, 2000; Livingstone & Fyfe, 2000; Werner and Dyer, 1994), studies of social and adaptive factors in communication (Noble, 1999; Wagner, 2000), models on the sensorimotor grounding of symbols and meaning (e.g. Cangelosi & Harnad, 2002; Steels and Kaplan, 1999; Kaplan, 2002), and computational studies on the neural and cognitive factors in language use, evolution, and origins (e.g. Arbib, 2002; Cangelosi & Parisi, 2001; Ellefson and Christiansen, 2000).

The majority of these models study the evolution of genetically-determined lexicons (e.g. Cangelosi & Parisi, 1998; Steels, 1996), or of languages that are only culturally transmitted (e.g. Kirby, 2001). Few studies simulate languages that depend on both genetic (i.e. evolution) and cultural (i.e. learning) factors (e.g. Cangelosi, 1999; Cangelosi & Harnad, 2002; Kirby & Hurford, 1997). However, these do not directly analyze the implications for phenomena such as the Baldwin effect. Other models directly address issues related to language learning and

evolution. Some study the critical period for language acquisition (e.g. Batali, 1994; Hurford, 1991). In Batali's (1994) simulation of language learning recurrent neural networks, it was shown that the connection weights undergo improvements through genetic assimilation. Just a few models of language evolution explicitly study the Baldwin evolution, and all focus on the assimilation of LAD based on the setting of principles and parameters (Briscoe, in press; Turkel, in press; Yamauchi, 2000). Yamauchi (2000) uses different genotype/phenotype mapping methods for evolving principle and parameter LADs. He adds pleiotropy to Turkel's (in press) genotype encoding model of LADs, meaning two or more genes express one feature of the phenotype based on a randomly generated look-up table. The results show that the Baldwinian evolution of the LAD is suppressed when the complex genotype/phenotype mapping is used.

In this paper we will extend Cangelosi's (1999; 2001) model on the emergence of compositionality to specifically investigate the Baldwin effect in the evolution of grounded languages. In particular, various experiments will be performed to investigate some of Mayley's (1996) ideas on the costs and benefits of learning with respect to the consequences they have on Baldwinian processes. In Cangelosi's original model a population of artificial agents first learn through evolution to forage successfully and then to communicate about food. Each agent learns to accomplish the linguistic task through learning (i.e. by supervised feedback on their neural network *during their lifetime*) and the foraging task through evolution (i.e. the population as a whole undergoes evolutionary pressure *across generations*). Adding these new modifications to Cangelosi's model will permit us to look at the interactions of both the learning and the evolution of language to see if there are any Baldwinian effects taking place, i.e. the assimilation of learned linguistic behaviors into the agent's genotype (Waddington, 1942). One way to investigate this is to record the populations' neural network pre-learning error before and after learning. If any Baldwinian processes have been at work, then it should be possible to see a decrease in the pre-learning error of the last generation, before any learning has taken place. Once a Baldwin effect has been achieved will be possible, following Mayley's (1996) suggestions, to manipulate certain aspects of the environment such that any Baldwinian effects can be increased or lessened. One possible way to implement this is to vary the rate of cultural variation in language production. The idea here is that if the language used between generations undergoes a large amount of variation across those generations, then there will be a strong inhibitory pressure against any Baldwinian effect taking place as each new generation cannot 'store' highly variable (and unreliable) properties of the

language. Similarly if the language undergoes little or no cultural variation and at the same time there exists costs in having to learn that language (e.g. by performing the foraging task poorly due to faulty learning) then there would be a strong pressure for the species to encode aspects of that language in the genome so that new individuals come into the environment in a prime position to take advantage of the linguistic structures already in place.

Finally, once the Baldwin Effect has been established, it is important to ask the question of just what is being assimilated into the genome of the agents in the population? Are the agents born into their world ready to speak the language of their ancestors without any lifetime experience or is it more subtle than that? Certainly - as has been discussed earlier - it is the case in humans that newborn children, if lacking contact with other language using humans, will not go on to develop their own language. Rather, it seems as though they come into the world with some innate ability to pick up the language that is being used in the surrounding environment. If there is no such language environment they cannot develop their own language. However, the ease with which they do pick up the language in use around them suggests some innate structures that are perhaps primed and ready to take in that language and, furthermore allows the child to make some amazingly correct educated guesses about the forms that the language will take. We hypothesize that one of the things that do become assimilated into the genome of a language using community of agents is a heightened sense of the boundaries between categories of environmental stimuli as well as a heightened sensitivity to the similarities between members of the same category. It is as if language helps to fill in the outlines of objects such that they stand out from their background in much the same way that drawing a thick line around an object in a picture helps to disembed it from all the other background clutter. Following on from Harnad (1987), this effect, known as Categorical Perception, will be investigated by recording the agents' hidden node activation at the beginning of stage 2 and again at the end of the simulation. A comparison will be made between the initial generation's hidden node activations and the last generations' hidden node activation to see if any categorical perception effects have taken place. This will show up as a contraction of the cluster of hidden node activations gained from similar objects and an increase in the distance between clusters gained from different kinds of objects (cf. Cangelosi & Harnad, 2002).

4. The model

This model is based on the grounded approach to studying language evolution (Cangelosi & Harnad, 2000; Parisi & Cangelosi, 2002). Specifically, the approach is to implement a simulated world in which simple agents learn to forage for foods, initially only through the process of evolution – meaning that each individual does not undergo any learning during its lifetime. This evolutionary learning will be developed through the use of a genetic algorithm operating on the connection weights of the agent's neural networks. Each agent will occupy a cell in a lattice world and will forage in an environment containing two types of food – edible and inedible “mushrooms”.

The overall setup of the model is identical to that used by Cangelosi (1999; 2001). Here we give a summary of the main features. An agent has to forage by identifying foods depending on the 18 bit representations of the visual features of mushrooms. In the environment there are 3 different categories of edible mushrooms (E1, E2, E3) and three types of poisonous mushrooms (P1, P2, P3). Four different exemplars of each category exist in each world. All edible mushrooms must be approached and eaten, and all toadstool must be avoided. Each sub-type of edible mushroom requires a specific action to be properly eaten, e.g. wash E1, cut E2, squash E3.

An agent will gain 1 energy point for any edible foods eaten but will have 1 energy point deducted for any poisonous (inedible) foods eaten. The agents' fitness is measured as the amount of energy left at the end of each epoch. The simulation on the evolution of foraging (stage 1) lasts for 300 generation. In the next 100 generations (stage 2) agents are allowed to communicate. Twenty expert foragers (i.e. the selected parents) are carried over into the next generation and act as speakers and teachers. Agents have to continue to successfully forage as before, but this time they have no access to the visual features of the food in 10% of the time. However, they always receive an additional linguistic input given to them by the 20 parents (each parent only speaks to its own offspring). During each cycle of interaction between the parent and its child, three different tasks are executed (Figure 1). In the first task, the parent sends out a signal that constitutes its own verbal description of the food that is the closest to the child. The child uses this linguistic information to decide what to do with the mushroom (only for 10% of the times can the child see the visual features of the mushroom). In task 2, the child agent will also perform a naming task in which it is shown the perceptual properties of the food

and it must produce a linguistic description of the mushroom. Through a backpropagation learning cycle, it compares and corrects its output with the teaching input provided by the parent. In the third task, the child performs a linguistic imitation task in which it uses as input only the parent's description, reproduce it as output, and once again backpropagation is applied.

4.1 The Neural Network

The behavior of an agent is controlled by a multi-layered feed-forward neural network (Figure 2). The network contains three sets of input units: three nodes to encode the position of the food, eighteen to encode the perceptual properties of the food, and eight linguistic input units. The hidden layer has five nodes. The output layer will comprise two groups of units: three encode the agents' actions (move in the environment and wash/cut/squash), and eight encode the linguistic output.

The linguistic input and output nodes are constructed of two winner-takes-all clusters, one of two nodes, and another of six nodes. Within each cluster, the node with the highest activation is set to 1 and the rest are set to 0. Each unit of a cluster encodes the activation of a "word". Thus at any time, two nodes/words are on, one per cluster.

In stage 2 of the experiment, the neural network of each agent will undergo supervised learning of the parents' verbal input. This will use the error backpropagation algorithm (Williams et al., 1986).

The connection weights used in the neural network are initially assigned randomly. They undergo changes through the genetic algorithm's stochastic operators and during error backpropagation in stage 2. Only the pre-learning weights are passed on to the next generation.

4.2 The Genetic Algorithm

A genetic algorithm is used to evolve the connection weights of the agents' neural networks. The weights are encoded in the genotype as a chromosome comprised of real numbers. The population consists of 80 agents. The genetic algorithm will select the best foraging agents

based on the fitness levels gained by eating the food sources during the agent's lifetime. The linguistic behavior does not produce any fitness gain directly. Once the top 20 agents have been selected, they reproduce asexually each making five offspring. The chromosomes of the new agents undergo random mutation whereby a normally distributed random number between ± 1 will be added to a randomly selected 10% of the genes. Here the search space is all the behaviors possible given the input to the networks.

During each generation, an agent lives in 20 different worlds (i.e. epochs) where foods are randomly placed.

4.3 Cultural Variation

As with human languages, where each new generation develops unique forms of expression and idiosyncratic meanings of already established words, the languages in this model are subject to the process of cultural variation which lets the languages change and evolve over generations. Cultural variation and language distortions occur at the interface between the old generation of language users and the new generation. At the end of each generation the output of the 20 most successful individuals' symbolic nodes are recorded and a random number between ± 0.5 is added to each linguistic output node before the winner-takes-all method is applied. This introduces the necessary variation in the language environment, which is needed for the experiments looking for a Baldwin effect.

4.4 Learning costs

Another aspect of the simulation that needs to be variable is the cost associated with learning the correct mapping between parental symbolic input and food sources. Mayley (1996) discusses the idea that one of the potential costs of learning is when a young agent, due to insufficient or faulty learning, performs a task wrongly or poorly and as a consequence fails to capitalize on some fitness benefit. In the present simulation, this can be viewed as the incorrect mapping of parental language into the right food category. For example a young agent may incorrectly interpret the symbol for a poisonous food as an edible one and proceed to eat the food source therefore incurring a penalty of -1 to its fitness score. Or alternatively an edible

food source may be avoided as the word given to represent the food was mistakenly translated as meaning a poisonous food source. For our purposes we will let the learning cost stand as the -1 fitness penalty for eating a poisonous food source. We can also reduce the learning cost to zero by removing this penalty, thus any poisonous food sources eaten can be made to have no effects whatsoever¹.

4.5 The Simulation Algorithm

Stage 1 (generations 1-300)

1. Create a population of 80 agents
 - a. Initialize their neural network connection weights
2. Place the population of agents in a series of 20 worlds (epochs), for 50 steps in each world.
 - a. At each time step the agent receives the angle of the nearest food and the 18 bits of its perceptual features
 - b. When the agent enters a cell where a food resides, it performs an action such as wash, cut or squash
 - c. Update fitness: If the food source is edible and the right action is performed: add 1 point to the agent's fitness score; If the food source is poisonous, subtract 1
4. After the 20 worlds, sort the agents on their fitness, pick out the top 20 agents, make five copies of genotype (= pre-learning weights) of the selected agent, mutate 10% of the weights of each new agent
5. Go to 2 until generation 300

Stage 2 (generations 301-400)

1. Keep the 20 parents from generation 300.
2. Place the 80 new organisms in a series of 20 worlds.

¹ There is also an additional cost associated in misinterpreting linguistic symbols associated to edible food categories. For example, no fitness point is gained when an E1 mushroom is treated as an E2 mushroom (e.g. cut instead of washed). Since this situation does not cause any fitness decrease, we prefer to manipulate only the cost associated to mistakenly approaching a poisonous food because of the negative (explicit) consequences in the fitness. The effects of fitness non-gains with edible mushroom is however consider later (see discussion of experiment 4 results in section 6).

- a. (Task 1: Action) At each time step the agent receives the angle of the orientation of the nearest food and a verbal description of that food which is initially a random two-symbol input. In 10% of the time the agent will also receive the food perceptual (visual) features.
 - b. Move the agent and update its fitness according to the learning cost (i.e. poisonous foods can be either give -1 or 0 fitness points)
 - c. (Task 2: Naming) Input only the perceptual properties of the food and apply back-propagation to the linguistic output units using the parent's words as teaching input
 - d. (Task 3: Imitation) Input only the two words provided by the parent and apply back-propagation to the linguistic output units using the same words as teaching input
 - e. Save the learning error for the first and last world
3. After the 20 worlds, expose each agent to all mushroom types and record the output lexicon. Apply noise to output nodes when cultural variation is on. Save the lexicon of the best 20 agents for the next generation.
 4. Sort the agents on their fitness, pick out the top 20 agents, make five copies of genotype (= pre-learning weights) of the selected agent, mutate 10% of the weights of each new agent
 5. Go to 2 until generation 400

5. Results

5.1 Initial Replication of Cangelosi (1999)

Initially we replicated Cangelosi's (1999) results in which a population of agents evolved both the ability to successfully forage and a compositional language (verb-noun languages). The majority of the parameters were exactly the same as in Cangelosi (1999). Only one parameter was significantly different. In the original model the cultural variation noise is always added for each learning trial. In the present model, noise is added once to the parents' lexicon. This permits consistency in the cultural transmission of language to all offspring of the same parent.

Ten baseline experiments with different randomly seeded populations were used. Eight populations managed to successfully forage for all edible food categories and avoid all poisonous food categories. Figure 3 shows the fitness levels gained for this first stage of the project. Stage 1 ran for 300 generations and figure 3 shows the fitness levels of the eight successful populations averaged together for both the best agent in the population and the top 20 best performing agents. By the last generation the best agents manage to collect on average 4.5 edible foods and avoid all poisonous foods. In the remaining two populations, some of the edible or poisonous category mushrooms are incorrectly dealt with by the agents.

The eight successful populations from stage 1 went on to the stage 2 simulations. Of these, seven successfully developed foraging strategies using both perceptual and linguistic information. The only population that failed to develop a successful strategy failed due to an inability to differentiate between edible food E2 and E3. Figure 4 shows the fitness for the seven successful populations. In comparison with stage 1, the difference in performance between the best agent and the top 20 agents is much closer. Furthermore, populations in stage 2 reach optimal foraging much quicker than the populations in stage 1. Stage 2 populations first reach a fitness level of 70 at around generation 90, whereas for stage 1 this kind of fitness is only reached around generation 150. Overall fitness by generation 300 is also higher in stage 2 than in stage 1, as in Cangelosi (1999), suggesting that the addition of the language component has beneficial effects on the agents' abilities to forage successfully. This could be due to an increased categorical perception effects in linguistic simulations, as outlined by Cangelosi & Harnad (2002). This issue will be explored further below.

The actual form that the evolved language took confirms Cangelosi's (1999) findings. Out of the seven populations that successfully foraged, one used the signals of only the 6-unit cluster to differentiate between the foods. This can be considered a single-word language. Two of the successful populations evolved combinatorial two-word pairings for the two major food types (i.e. edible/approach and poisonous/avoid) and also developed signals to differentiate between the three different subtypes of edible food types. However, there was no compositionality since signals were not systematically associated to identifiable meaning units (e.g. no signal was consistently associated to "avoid" or to "approach"). The final four populations developed a robust compositional communication system based on verb-noun languages. Most of the population shared the same symbols ("verbs") for the approach/avoid

actions, and other symbols (“nouns”) were systematically linked to the different kinds of food types E1, E2 and E3.

Figure 5 shows one of these verb-noun languages. The two columns of circles (to the right) represent the two different types of verbs available to the agents. The symbols E1, E2, E3 and P1, P2, and P3 refer respectively to the three types of edible foods and the three types of poisonous foods. The block of six columns on the left corresponds to the 6-node winner-takes-all cluster. This potentially contains enough units for the agents to come to identify each type of food source with a unique node (i.e. single-word language). The block of 2 units on the right correspond to the smaller 2-node cluster. The height of the cone corresponds to the probability of use of each node: a full cone corresponds to the 100% use of the node in the whole population. In this figure, E1, E2 and E3 each receive the same symbol in the 2-node cluster, whilst each subtype of edible food (E1, E2 and E3) is allocated a unique node in the cluster of six nodes. This is necessary for the agent to differentiate between the three subtypes. The poisonous food sources again all receive the same node activation in the 2-node cluster, though critically this is different from the one assigned to the edible food sources. As there is no need for the agents to differentiate between the different subtypes of poisonous foods - as the response required is simply to avoid - the poisonous food sources are all allocated the same node activation in the 6-node block.

The model will now be used for new simulations where the learning cost and the cultural variation parameters are systematically varied to study their consequences on the Baldwin effect.

5.2 Experiment with the Baldwin Effect

Mayley (1996) argues that the Baldwin effect will be most prominent when the costs of learning outweigh the benefits. In the present model we expect to see a Baldwinian effect when these three criteria are met:

1. Agents undergo both evolutionary adaptation and lifetime learning.

2. There are predictable regularities in the environment, with no cultural variation²
3. The costs of learning (i.e. negative fitness for poisonous foods) outweigh the benefits

With these criteria in mind, it is possible to pick out the two specific components of the simulation setting that we have explained to be expected to affect the appearance and amount of Baldwinian processes, namely the amount of cultural variation (0 or ± 0.5) inherent in the passing on of languages to the next generation and the costs of learning (0 or -1 for poisonous foods). Next we present the results gained from systematically manipulating both of these factors. To assess the presence and the amount of Baldwin effects, we measure the Root Mean Square (RMS) error of the best twenty neural networks in all epochs of the first generation of stage 2 (gen. 301) and in all epochs of the last generation (gen. 400). The genotypes of the eight successful populations of stage 1 (i.e. end of generation 300) will be used for running the Baldwin effect simulations.

5.2.1 Experiment 1 Baldwin Effects on RMS scores with cultural variation set at ± 0.5 and learning costs set to -1.

In this first experiment we looked at the extent of Baldwinian processes with the standard settings similar to that of Cangelosi (1999). Cultural variation was set at 0.5 which means that, at the time that the language is passed from one generation to the next, a random amount of ± 0.5 is added to the linguistic output nodes of the parents before the winner-takes-all method computes the only active node in a cluster. This had the effect of a probabilistically changing the form of the language by a small random amount, similar to what happens in real human languages. In addition, the costs associated with the use of an erroneous language (e.g. when poisonous foods are eaten by mistake) is to decrease the fitness score by 1³. Figure 6 shows the RMS error for the initial generation (301) through the agents' lifespan in all twenty epochs and the same for the last generation (400).

² Though this in itself does not guarantee a Baldwin effect, as either these regularities will be left untracked by evolutionary processes (e.g. when the cost in learning them throughout the agent's lifetime is negligible), or when the cost becomes prohibitive.

³ Note that at this stage in the experiment all organisms have already developed the ability to discriminate between all foods without error based on their perceptual appearance. Thus any mistakes made after the introduction of language in stage 2 reflects the agents' difficulties in mapping the language onto the previously learnt categories.

With these settings, agents in the last generation actually begin their lifespan with a slightly higher RMS (0.148) error than the initial generation (0.143), with a difference of only 0.005. However, as can be clearly seen in Figure 6, the RMS error quickly drops for the last generation, resulting in a large deviation in RMS scores in later epochs. Using as independent variable the evolutionary amount of species learning (with two levels: first and last generation) and as dependent variable the RMS error, a t-test independent measures design was used to compare the average error of the first and last generations. This test resulted in significant difference of the two means with a value of $t = 435.03$ and $p < 0.0005$. These results suggest that although little, or none, about the specifics of the task (i.e. the actual lexicon) has been initially canalized, some ability to subsequently learn the linguistic task efficiently has been assimilated, making agents more effective language learners. This will be discussed in more detail in section 6.

5.2.2 Experiment 2. Baldwin Effects on RMS when Cultural Variation = 0 and learning costs set to -1.

By reducing the cultural variation it is hypothesized that Baldwinian effects will increase. As the environment becomes fixed and thus more predictable (in these experiments the term environment refers to the linguistic stimuli), genomic structures might become able to predict what the agent will be presented with at birth. As a consequence, features of the linguistic environment can be assimilated in the genotype to anticipate the environment.

Figure 7 shows the RMS error for both start and end populations with no cultural variation and a learning cost of -1. The difference between the mean RMS errors of the two generations is significant at $t = 16.78$ and $p < 0.0005$. In addition, the end generation clearly has a much-increased innate capacity to perform the task since the first epoch. The initial RMS error rate (0.155) is already very low when compared to the starting RMS of the initial generation (0.067).

The comparison of the figures of experiments 1 and 2 also shows that the slope of learning between the two generations is parallel in Figure 7. Instead, in Figure 6 the end generation has a much steeper drop in RMS, especially in the first eight epochs, after which it levels off near the optimal performance value of zero. This could be a reflection of the different types of characters internalized by Baldwinian processes in the two learning environments. The

environment in experiment 1 was one where variability was common and evolution found ways to evolve innate characters that made learning efficient once the nature of the language environment was presented to the agents. Experiment 2 contained no such variability in the language environment and this enabled the evolutionary process to encode aspects of the forthcoming language into the genome. This results in agents that were quickly up to speed in being able to use that given language and little subsequent need for learning past the first few epochs.

5.2.3 Experiment 3. Cultural variation set to ± 0.5 and no learning costs

In this environment the agents must remain flexible in their initial genomic structure if they are to learn the varying language from generation to generation. However, as the costs of learning the languages are nil – i.e. any errors made have no consequence in terms of fitness – there should be no pressure to incorporate the ability to learn languages into the pre-lifetime learning genome.

Figure 8 shows that, although there has been some assimilation, it does not reach the extent that was seen for when cultural variation was set to zero. The initial starting point for both generations differs, and the rate of learning thereafter remains similar. So it seems that even though cultural variation is present, the rate of variation (± 0.5) was low enough to allow some aspects of it to be assimilated into the genome. However, the lack of any learning cost meant that agents were not much pressured to assimilate learning properties between the first and last generation, as shown by the similar rate of learning. The difference between the two conditions was significant at $t = 9.936$ and $p < 0.0005$, with a difference between the initial RMS scores of both generations of 0.032.

5.2.4 Experiment 4. No cultural variation and no learning costs

Agents in this type of environment are expected to show an increased level of assimilation of the features of the language, because it stays stable over all generations. However, as there is no cost associated with learning the language, the amount of assimilation should be somewhat less than in experiment 2 where the lack of cultural variation was combined with learning costs. Figure 9 shows this to be the case. The last generation starts out from a point in the

learning space much closer to the goal state than the first generation. The distance between the two is slightly less (a difference of 0.056) than between the two generations in experiment 2 (a distance of 0.088). The overall difference between the two conditions is significant at $t = 16.9832$ and $p < 0.0005$.

6. Discussion

The main findings of this simulation are here summarized:

1. Confirmation of Cangelosi's (1999) findings that a population of interacting agents is capable of evolving a shared communication system based on a compositional (e.g. verb-noun) language.
2. The manipulation of the cultural variation of language between generations makes it possible to produce Baldwinian effects such that aspects of the language can become innate.
3. Similarly, by varying the cost of learning a language Baldwinian processes can be induced.
4. Depending on which variables in the language learning environment are manipulated, the precise nature of the linguistic abilities assimilated into the genome changes.
5. By fixing the structure of the language across generations, aspects of the specific structure of the language can be assimilated.
6. However, if the structure of the language is allowed to vary by cultural variation, Baldwinian processes cause the assimilation of a predisposition to learn, rather than assimilating any structural properties associated with a specific language.

Regarding finding 1, it should be pointed out that whilst the results of this project corroborate Cangelosi's (1999) findings, there are some differences in the new simulations. One of the main differences in terms of results between these simulations and Cangelosi's was in the difference in performance in stage 2. In Cangelosi (1999) the addition of a linguistic component to the agents' world produced an acceleration in their ability to learn how to forage that is much less pronounced in the simulations here. However, the effect is still present, though somewhat reduced, and this could merely be down to the differences in the details of implementation. Both sets of simulations show an increased performance overall as reflected in the higher levels of fitness at the end of stage 2 compared to stage 1 (see figures 3 and 4 in the

results section). What could account for this increase in the speed of learning and the overall performance? Cangelosi and Harnad (2000) report a study that shows how the addition of language can help to improve categorical perception effects (Harnad, 1987), i.e. to compress within category distances, whilst at the same time increasing between category distances. During training, a neural network will alter its weights in such a way to partition the search space such that similar patterns of inputs are represented by similar hidden unit activations; conversely, dissimilar input activations will come to be represented by dissimilar hidden unit activations. By adding a symbolic (linguistic) component to an already learned pattern recognition task, Cangelosi, Greco & Harnad (2000) showed how the hidden unit activations become even more similar for members of the same category and less similar between members of different categories. This has also been observed in languages based on verbs and nouns (Cangelosi & Parisi, 2001).

In the current set of simulations, the addition of language could have increased categorical perception effects and help the agents to better distinguish between food categories. This would increase their ability to perform the foraging task. To test this hypothesis, we measured the agents' hidden node activations for all the different types of edible mushrooms both before and after the addition of language. Though the agents performed well in distinguishing between the different types of edible foods at the end of stage 1, we wanted to see if their increased performance at the end of stage 2 was due to a more differentiated set of hidden unit activations for each sub-category of edible foods. Categorical perception effects can be measured by comparing the Euclidean distances in the multidimensional space of the hidden node activation. Any particular set of N hidden node activations produces a unique coordinate in an N -dimensional space (5 dimensions in this model with 5 hidden nodes). Comparing the hidden node activations to the different exemplars of each type of edible food produces a cluster of co-ordinates. Any food that is strongly differentiated from other foods will produce a small cluster of activations that in turn will be separated from other clusters representing different categories of foods. Thus an improvement in discriminatory performance will show up as a decrease in cluster size and an increase in distance between other clusters – the within and between category effects mentioned above. Figure 10 shows the difference in edible food cluster size for each type of edible food both before the addition of language (init) and after (end). For each type of edible food (E1, E2 and E3) we can observe a drop in the size of the cluster as we move from the initial generation without language to the end generation with language. Although the difference is minimal for E1, this is evident for E2 and E3 and it shows

that at least some edible foods have become more homogenous through the use of shared linguistic signals.

In addition, each type of edible food has become more dissimilar to the others thanks to the addition of a linguistic marker as shown in figure 11. The dark columns represent the average distance between each of the food types before the addition of language and the light columns the distance after the addition of language. The difference in height between the dark and light columns shows the increase in distance between each food subtype. Therefore this confirms that by evolving a shared language the agent population has decreased its ability to recognize variations within the food categories, as shown by the smaller clusters, and increased its ability to recognize variations between food categories as shown by the increased distance between clusters.

Figure 12 shows the distance between each edible sub-category and the inedible/poisonous foods. The first thing to notice is that even the initial distances between the edible and inedible foods are larger than the distances between the different sub-categories of edible foods as shown in figure 11. The smallest distance between an edible food source and the inedible foods is about 1, nearly the same as the largest distance between any of the edible food sources. This is due to the higher level of difference in the bit representations between edible and inedible and to the different motor responses associated to each food category. Thus even in the initial, pre-language population the difference between the edible and inedible categories is high. Again, though the post-linguistic population has much more separation between edible and inedible clusters (c.f. increased height of the light columns of Figure 11), similarly the inedible cluster (of which there is only one as the agents did not have to differentiate in terms of behavior between the different sub-types of inedible foods) has also undergone shrinkage in size as shown in Figure 13.

Overall, this categorical perception effects show that the addition of linguistic markers for both the objects themselves and the associated actions help the agents to discriminate between different categories and subcategories (see also Cangelosi & Harnad, 2000). As a consequence, the behavior of foraging becomes more efficient and less error prone as shown by the rate of fitness increase and the overall heightened fitness levels at the end of stage 2 (Figure 5).

This simulation also shows interesting phenomena on the effects of cultural variation and learning costs. For humans, the language environment is constantly changing. Each generation recreates only to some extent the language of its forefathers. One has only to look back to Chaucer or Shakespeare to notice how much the English language has changed over the generations. In the first experiment of the model, this was reflected in the cultural variation set at ± 0.5 at each generation. Furthermore, it is trivial to imagine situations in which an individual with a faulty understanding of the language in use around him could suffer some injury either to his physical well-being or perhaps more subtly to his social standing. This too is reflected in experiment 1 by the fitness costs of mistakenly eating a poisonous food. Prior to the introduction of language, all agents had proven to be good foragers in which categorization mistakes were minimal. Any subsequent failings in performance are therefore due to a faulty understanding of the language in use, or to the use of bad languages. Thus experiment 1 most closely reflects the situation in human language environments. Looking at the results of this experiment (Figure 6) we see that there appears to be no benefits to the last population in terms of initial ability to reproduce the language of the previous generation. In fact, initial RMS error scores are slightly higher as shown below in figure 14 (first column pair *cv lc*). This is to be expected when one considers the fact that as the language changes every generation there is little stable linguistic information to be learnt and assimilated into the genome of the agents. This echoes Deacons (1997) comments on the effects of language change:

“The relative slowness of evolutionary genetic change compared to language change guarantees that only the most invariant and general features of language will persist long enough to contribute any significant consistent effect on long-term brain evolution.” (Deacon 1997: 329)

Although we cannot directly compare the extremely simple forms of languages developed by the agents with real human languages, certain interesting similarities are present. Human infants come into the world with what Greenough et. al. (1987) would call “Experience Expectant” (EE) systems that makes them ready to absorb all the correct details of the language environment they find themselves in. This seems to be a strong account for the extreme speed and efficiency with which children become language experts. What the agents possess is, like human infants, a predisposition to efficiently learn the current language. This can be seen in the difference in the rate of learning between the initial, language naïve generation and the last generation. The first column pair in figure 15 shows this difference. Agents in the last

population are much quicker to pick up the language than those agents in the first generation. Some features related to language learning abilities have managed to find their way into the innate characters of later generations.

We can also compare the learning rate for experiment 1 with experiment 3 in which the cost of learning is reduced to 0. Though there is still some increase in the speed of learning for the last generation, the effect is much lessened (column pair $cv \neg lc$ in Figure 15). When there are no costs associated with learning, the pressure to assimilate the aspects of the task into the genome in order to produce more efficient behavior is lessened. The overall fitness of the agent at the end of its lifetime is the only measure available to evolutionary processes and if the behavior of the agent produces no changes in this overall fitness measure, then evolution ignores that behavior. Instead, in experiment 2 where cultural variation has been removed but learning costs are still present, there is a difference between the pre-learning performance of agents in the initial and final generations (column pair $\neg cv lc$ in figure 14). Here, due to the unchanging nature of the language environment and the impetus of Baldwinian processes brought to bear by learning costs, we can see that the initial performance of the last generation of agents far outstrips the first generation. It seems that in the last generation, agents come into the world already speaking the specific language. Whilst this is obviously not the case for human infants' language, it is possible to think of other abilities, or structures, that have similarly become innate (Elman et al., 1986). In Experiment 2, the gain reflected in the initial performance level is offset by the accompanying drop in the overall learning rate as shown by the second column pair $\neg cv lc$ in figure 15. As most of the characteristics of the language have been assimilated into the pre-learning genome of the agents, there is little left to learn after birth. Thus the specification of a learning system, an EE system in Greenough's terms, becomes less important.

Finally, experiment 4 presented the agents with a stationary language environment with no associated learning costs. Though agents still assimilate aspects of the task into the genome (see column pair $\neg cv \neg lc$ in Figure 14), the level of assimilation is not so high as in the other case with no cultural variation. The explicit learning cost of a negative impact on fitness levels was removed from the experiment, but there still exists subtle learning costs if the agents refuse to eat edible foods available due to faulty language understanding. This indirect fitness cost could explain the small assimilation effect seen here. Likewise, there is a small increase in learning efficiency for the last generations as shown by the fourth column pair of figure 15.

7. Conclusions

The analysis of the present model clearly shows that there are different types of Baldwin effects in language evolution. If the learning environment is changing and there are associated fitness costs in learning the task, then one strategy the population can adopt is to produce individuals that are quicker at learning the language at hand. For this to happen, it is necessary that, even though the task is changing at the superficial level, there must be some underlying and singular method for learning that task. In the case of our agents with simple languages, the task mainly consists of mapping a new set of sensory inputs (the language) onto an already formed set of grounded internal representations. This task, though extremely simplified, still shares a basic trait of human languages, i.e. the use of symbolic relationships. As Deacon (1997) points out, for a communication system to qualify as a language, it must also develop logical relationships between the linguistic symbols themselves. The language used by the agents is compositional and uses syntactic relationships, such as that between verbs and nouns (Cangelosi, 2001).

The second way in which Baldwinian processes can be seen to work in the model is perhaps the more usual kind of assimilation, where actual specific behaviors become canalized into the genome. This occurs when the particular learning environment becomes fixed, thus enabling the details of the linguistic behavior to be assimilated. In our experiments, the lack of any cultural variation meant that even the structure of the language could be built into the agent's pre-learning genome. Whilst unrealistic in terms of human language, this does reflect other kinds of behavioral canalization.

Overall, this model support Deacon's (1997) ideas on the inheritance of general underlying capabilities that help to enable new generations to learn the linguistic system in place. Nothing in the experiments points to anything as complicated and specific as a Chomskian language acquisition device (LAD). In addition, no evidence from the research is available to support Pinker's (1994) claim that certain principles and parameters of the language are assimilated by Baldwinian processes. Rather, the behavior of the agents seem to reflect, in the case of culturally varied language, a general ability to more efficiently learn whatever linguistic structures are presented to the agents. Moreover, the analyses of the categorical perception

effects induced by language suggest that what is assimilated through Baldwinian processes, is a fundamental cognitive ability at the base of cognitive and linguistic behavior, that is building discrete categories of the objects in the world (Harnad, 1987; 1996). In the agents, this is realized through the evolution of neural networks whose initial weight setups facilitate language learning. This is consistent with Elman et al. (1996) “rethinking” of innateness.

Further investigations are needed to ascertain what precise mechanisms are at work in those agents with a pre-disposition to learn the language they are presented with. Analyzing neural networks is a non-trivial undertaking, but it would perhaps be interesting to enable the networks to switch on or off their level of plasticity in response to the amount of cultural variation within the environment. This would enable us to build a plasticity profile within the architecture of the network. In addition, it would also be interesting to experiment with different neural network architectures to see if such properties could have an effect in the efficiency of developing symbolic relationships between words and objects, as well as between the words themselves.

The cultural variation used in these experiments differed slightly to that used in Cangelosi (1999) where the variation was added not once at the interface between each generation, but over the many interactions between parent and child agents as the language was taught. Cangelosi’s method may prove to be more natural and comparing cultural variation themes may help in further understanding its relationship with the type and amount of Baldwinian effects on show.

Of course other aspects of the agents’ environment are open to variation. It would be possible to slowly alter the appearance of the food sources within a population. This, combined with various levels of linguistic variation, could act in synergistic ways and would be a rich source of experimentation.

Finally, in developing such models of the evolution of language, it is important to bear in mind the necessity of incorporating more of the details inherent in human languages and available experimental data (Tomasello, 2002).

Acknowledgements

Angelo Cangelosi has been partially supported by the UK Engineering and Physical Sciences Research Council (EPSRC Grant: GR/N01118).

References

- Ackley D., Littman M. (1991). Interaction between learning and evolution. In Langton C. Taylor C. Farmer, J. & Rasmussen S. (Eds.), *Artificial Life II*, pp.487-509.
- Arbib M.A. (2001). Grounding the mirror system hypothesis for the evolution of the language-ready brain. In Cangelosi A., Parisi D. (Eds.), *Simulating the evolution of language*, London: Springer
- Arita T, Taylor CE (1996). A simple model for the evolution of communication. In: Fogel LJ, Angeline PJ, Bäck T (eds) *Evolutionary Programming V*, MIT Press, Cambridge
- Baldwin J.M. (1896). A new factor in evolution. *The American Naturalist*, 30.
- Batali J. (1994). Innate biases and critical periods: combining evolution and learning in the acquisition of syntax. In Brooks R., Maes P. (eds), *Proceedings of the Fourth Artificial Life Workshop*. MIT Press, Cambridge MA
- Batali J. (1998) Computational simulations of the emergence of grammar. In: Hurford J., Knight C., Studdert-Kennedy M. (eds) *Approaches to the evolution of human language: Social and cognitive basis*. Cambridge University Press, Cambridge UK, 405-426
- Belew R.K., Mitchell M. (eds) (1996). *Adaptive individuals in evolving populations: Models and algorithms*, Massachusetts: Addison-Wesley
- Brighton H. (2002). Compositional syntax from cultural transmission. *Artificial Life*, 8(1): 25-54
- Briscoe E.J. (2000). Grammatical acquisition: Inductive bias and coevolution of language and the language acquisition device. *Language*, 76: 245-296
- Briscoe, E.J. (in press). Grammatical Assimilation. In M. Christiansen & S. Kirby (eds), *Language evolution: the states of the art*. Oxford University Press.
- Cangelosi A. (1999). Modeling the evolution of communication: From stimulus associations to grounded symbolic associations. In D. Floreano, J. Nicoud, F. Mondada (Eds.), *Advances in Artificial Life (Proceedings ECAL99 European Conference on Artificial Life)*, Berlin: Springer-Verlag, 654-663

- Cangelosi A. (2001) Evolution of communication and language using signals, symbols and words. *IEEE Transactions in Evolutionary Computation*, 5(2): 93-101
- Cangelosi A., Greco A., Harnad S. (2000). From robotic toil to symbolic theft: Grounding transfer from entry-level to higher-level categories. *Connection Science*, 12: 143-162
- Cangelosi A., Harnad S. (2000) The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution of Communication*, 4(1): 117-142
- Cangelosi A., Parisi D. (1998) The emergence of a 'language' in an evolving population of neural networks. *Connection Science*, 10: 83-97
- Cangelosi A., Parisi D. (2001) How nouns and verbs differentially affect the behavior of artificial organisms. In Moore JD, Stenning K (eds), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum Associates, pp 170-175
- Cangelosi A., Parisi D. (Eds.) (2002). *Simulating the evolution of language*. London: Springer-Verlag
- Chomsky N. (1975). *Reflections on language*. Pantheon Press.
- Chomsky N. (1991). Linguistics and cognitive science. In Kasher, Asa (eds.), *The Chomskyan Turn: Generative Linguistics, Philosophy, Mathematics, and Psychology*. Blackwell.
- de Boer B. (2000). Self organization in vowel systems. *Journal of Phonetics*, 28: 441-465
- Deacon T. (1996). Prefrontal cortex and symbol learning: Why a brain capable of language evolved only once. In B.M. Velichkovsky, D.M. Rumbaugh (Eds), *Communicating meaning: The evolution and development of language*. Mahwah NJ: LEA Publishers.
- Deacon, T. (1997). *The Symbolic Species: the Coevolution of language and human brain*. London: Penguin.
- Ellefsen M.R., Christiansen M.H. (2000). Subjacency constraints without universal grammar: Evidence from artificial language learning and connectionist modeling. In *The Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, Erlbaum, Mahwah NJ, pp 645-650
- Elman J.L., Bates E.A., Johnson M.H., Karmiloff-Smith A., Parisi D., Plunkett K. (1998) *Rethinking innateness*. MIT Press: Cambridge Massachusetts.
- Fodor J. (1983) *The Modularity of mind*. Cambridge, Mass: MIT Press.
- Gazzaniga M.S., Ivry R.B., Mangun G.R. (1998) *Cognitive neuroscience: The biology of the mind*. New York: Norton.

- Greenough W.T., Black J.E., Wallace C.S. (1987). Experience and brain development. In Johnson M.H. (Ed.), *Brain development and cognition: A reader*. Oxford: Blackwell.
- Hare M., Elman J.L. (1995) Learning and morphological change. *Cognition*, 56: 61-98
- Harnad S. (Eds) (1987) *Categorical Perception: The groundwork of cognition*. New York: Cambridge University Press.
- Harnad S. (1996) The Origin of Words: A Psychophysical Hypothesis In Durham, W & Velichkovsky B. (Eds.), *Communicating Meaning: Evolution and Development of Language*. NJ: Erlbaum.
- Hashimoto T. (1997) Usage-based structuralization of relationships between words. In P. Husband, I. Harvey (eds), *Proceedings of the Fourth European Conference on Artificial Life*. MIT Press, Cambridge MA, pp 483-492
- Hauser M.D. (1996). *The evolution of communication*. MIT Press, Cambridge MA
- Hazlehurst B., Hutchins E. (1998) The emergence of propositions from the co-ordination of talk and action in a shared world. *Language and Cognitive Processes*, 13(2/3): 373-424.
- Hebb D. (1949) *The organization of behavior: A neuropsychological theory*. New York: John Wiley & Sons.
- Hinton G.C., Nowlan S.J. (1987). How learning can guide evolution. *Complex Systems*, 1
- Hurford J., Studdert-Kennedy M., Knight C. (eds) (1998). *Approaches to the evolution of language*. Cambridge University Press, Cambridge UK
- Hurford J.R. (1991). The evolution of the critical period for language acquisition. *Cognition*, 40: 159-201.
- Kirby S. (1999). *Function, selection and innateness: The emergence of language universals*. Oxford: Oxford University Press
- Kirby S. (2001). Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2): 102-110
- Kirby S. (in press). Natutal language and artificial life. *Artificial Life*
- Kirby S, Hurford JR (1997) Learning, culture and evolution in the origin of linguistic constraints. In P. Husband, I. Harvey (eds), *Proceedings of the Fourth European Conference on Artificial Life*. MIT Press, Cambridge MA, pp 493-502
- Livingstone D., Fyfe C. (2000). Modeling language-physiology coevolution. In Knight C, Studdert-Kennedy M, Hurford J (eds), *The evolutionary emergence of language: Social function and the origins of linguistic form*. Cambridge University Press, pp 199-215

- MacLennan B.J., Burghardt G.M. (1994) Synthetic Ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2: 151-188
- Mayley G. (1996). The evolutionary cost of learning. In Maes P., Mataric M., Meyer J-A., Pollack J., Wilson S. (Ed.), *From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour*, MIT Press.
- Noble J. (1999). Cooperation, conflict and the evolution of communication. *Adaptive Behavior*, 7(3/4): 349-370
- Nolfi S. (2000). How learning and evolution interact: The case of a learning task which differs from the evolutionary task. *Adaptive Behavior*, (7) 2:231-236
- Nolfi S., Floreano D. (1999). Learning and evolution. *Autonomous Robots*, 7(1): 89-113
- Nolfi S., Elman J.L. & Parisi D. (1994). Learning and evolution in neural networks. *Adaptive Behavior*, (3) 1:5-28.
- Nowak MA, Komarova NL (2001) Towards an evolutionary theory of language. *Trends in Cognitive Science*, 5(7): 288-295
- Nowak MA, Plotkin JB, Jansen VAA (2000) The evolution of syntactic communication. *Nature*, 404: 495-498
- Oliphant M. (1999). The learning barrier: Moving from innate to learned systems of communication. *Adaptive Behavior*, 7: 371-383
- Parisi D. (1997). An Artificial Life approach to language. *Mind and Language*, 59: 121-146.
- Parisi D., Cangelosi A. (2002). A unified simulation scenario for language development, evolution, and historical change. In Cangelosi A., Parisi D. (Eds.), *Simulating the evolution of language*, London: Springer
- Pinker S., Bloom P. (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, 13: 707-784
- Pinker S. (1994) *The Language instinct: the new science of language and mind*. Allen Lane.
- Skinner B.F. (1953). *Science and human behavior*. Colliler-Macmillan: New York
- Steels L. (1996). Self-organizing vocabularies. In Langton CG (Ed.), *Proceedings of Artificial Life V*. Nara
- Steels L. (1997). The synthetic modelling of language origins. *Evolution of Communication*, 1: 1-34
- Steels L., Kaplan F. (1998). Situated grounded word semantics. In: *Proceedings of IJCAI-99*. Morgan Kaufman Publishing, Los Angeles, pp 862-867
- Steels L., Kaplan F. (2000). AIBO's first words: The social learning of language and meaning. *Evolution of Communication*, 4(1)

- Teal T.K., Taylor C.E. (2000). Effects of compression on language evolution. *Artificial Life*, 6: 129-143
- Tomasello M. (2002). Some facts about primate (including human) communication and social learning. In Cangelosi A, Parisi D (eds), *Simulating the evolution of language*, London: Springer
- Tomasello M., & Brook P.J. (1999). Early syntactic development: A Construction Grammar approach. In M. Barrett (Ed.), *The Development of Language*, Philadelphia, PA: Psychology Press, pp. 161-190
- Turkel W.J. (2002). The learning guided evolution of natural language. In T Briscoe (Ed.). *Linguistic evolution through language acquisition: Formal and computational models*. New York: Cambridge.
- Turney P., Whitley D., Anderson R.W. (1996). Evolution, learning, and instinct: 100 years of the Baldwin effect. *Evolutionary Computation*, 4(3): iv-viii. (special issue on the Baldwin effect)
- Waddington C.H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150
- Wagner K. (2000). Cooperative strategies and the evolution of communication. *Artificial Life*, 6(2): 149-179
- Werner G.M., Dyer M.G. (1994). BioLand: a massively parallel environment for evolving distributed forms of intelligent behavior. In Kitano H (Ed.) *Massively parallel artificial intelligence*. MIT Press, Cambridge MA
- Yamauchi Y. (2000). Can the Baldwin effect really explain the evolution of the LAD? In J.L. Desalles, L. Ghadakpour (eds), *Proceedings of the 3rd International Conference on the Evolution of Language*, ENST Paris, 263-265

Figures




Figure 1 – Learning tasks in stage 2

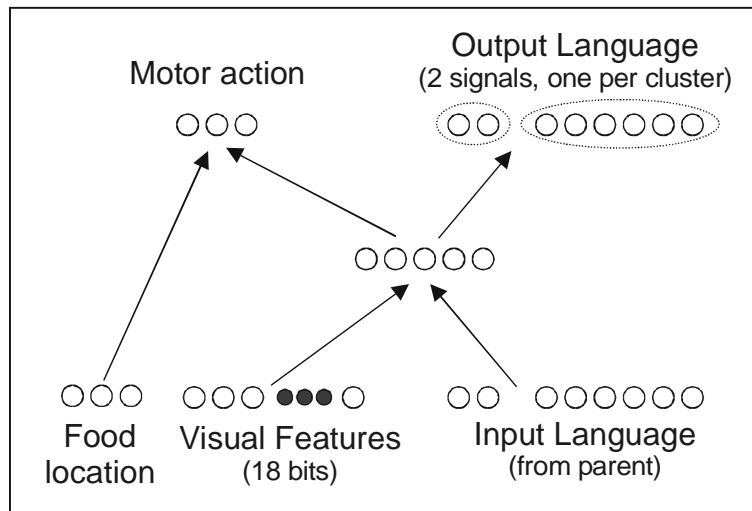


Figure 2 – Network architecture




Figure 3 – Stage 1 Fitness

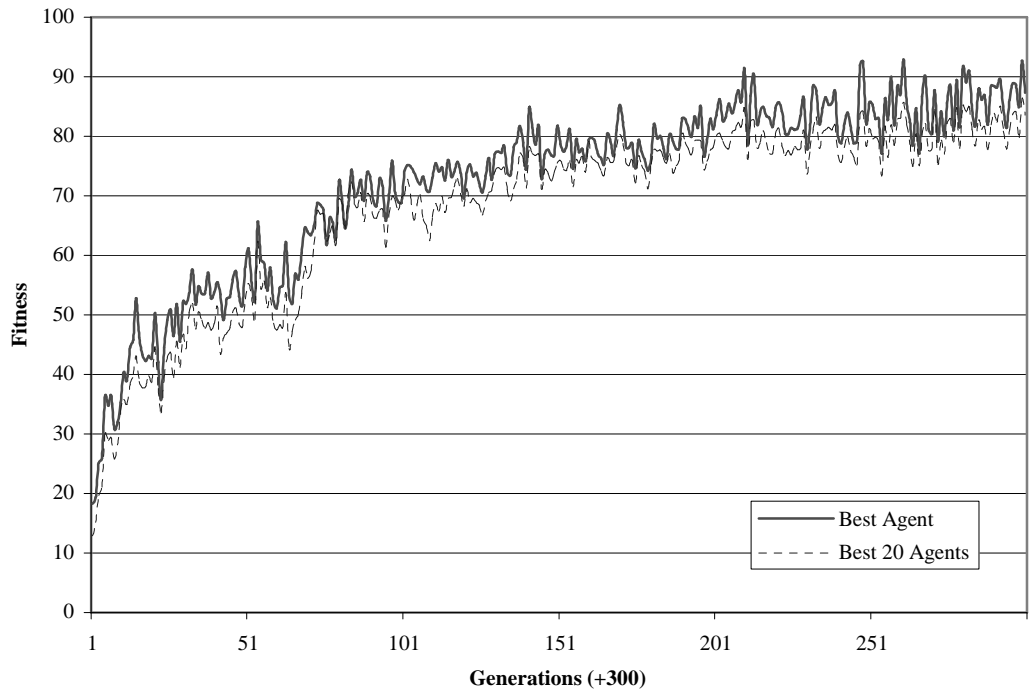


Figure 4 Stage 2 Fitness




Figure 5 – Verb-noun lexicon




Figure 6 – RMS Error for first and last generations with cultural variation and learning costs




Figure 7 – RMS Error with no Cultural Variation and Learning Costs




Figure 8 – RMS Error with cultural variation and no learning costs




Figure 9 – RMS error with no Cultural Variation and no Learning costs

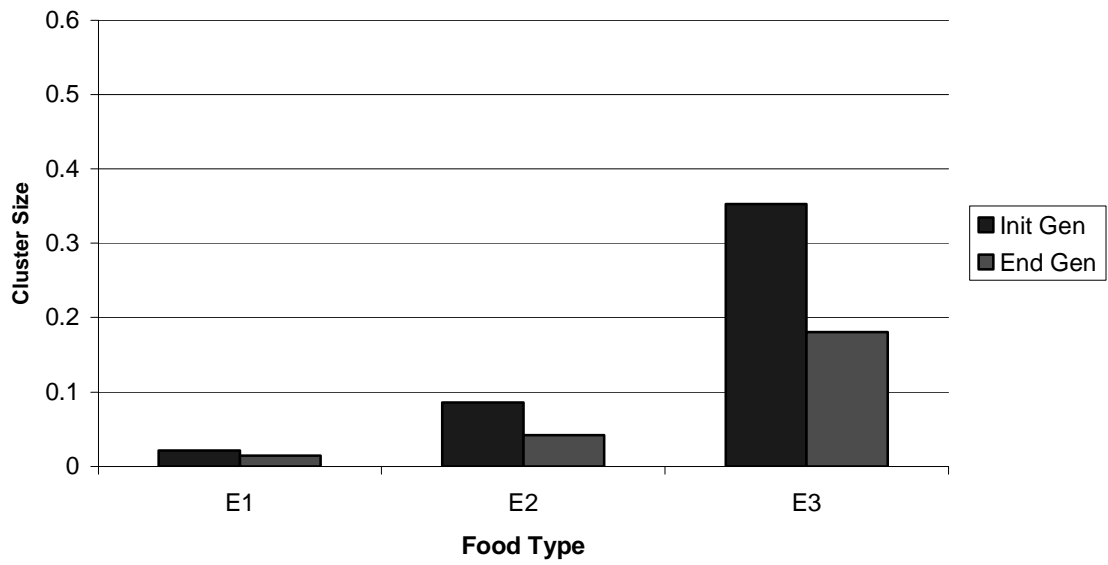


Figure 10 – Edible cluster size for initial and end generations

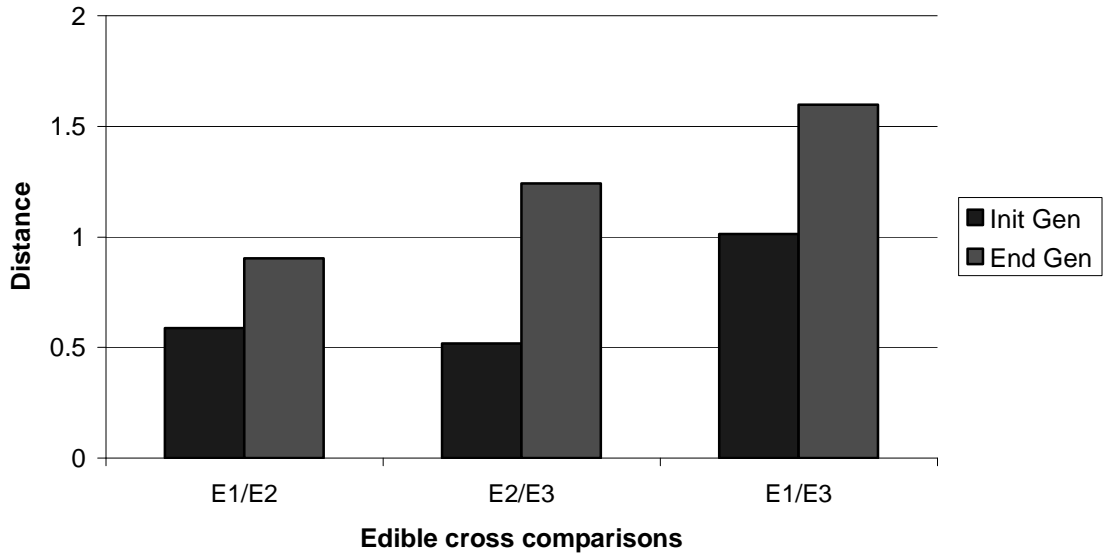


Figure 11 – Distance between edible clusters

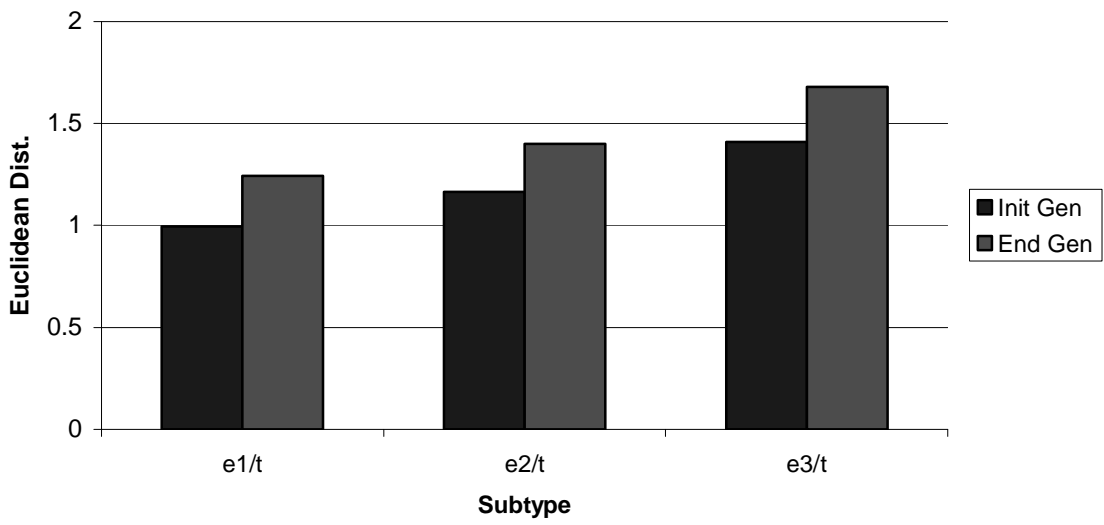


Figure 12 – Distance between Edible subtypes and inedible

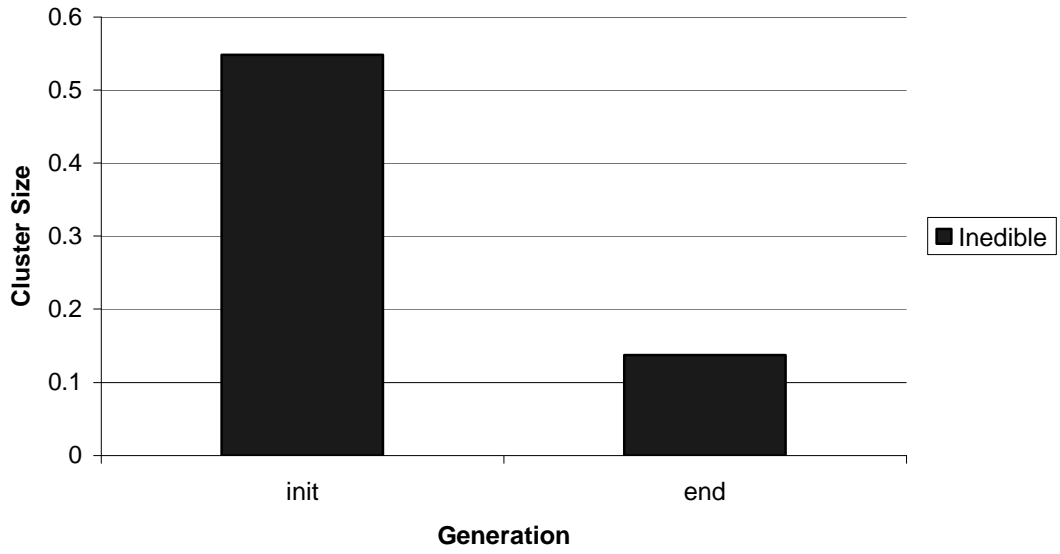


Figure 13 – Inedible cluster size




Figure 14 – Pre-learning difference of RMS in epoch 1 for initial and end generations in the four experiments.




Figure 15 – RMS difference in learning efficiency between the first and the last epochs.