Learners are losers: Natural selection and learning in the evolution of communication

Kenneth Smith
Language Evolution and Computation Research Unit
Department of Theoretical and Applied Linguistics
University of Edinburgh

kenny@ling.ed.ac.uk
Phone: +44 131 650 6657
Fax: +44 131 650 3962

1ST DRAFT OF PAPER SUBMITTED TO THE PROCEEDINGS OF THE 3RD CONFERENCE ON THE EVOLUTION OF LANGUAGE

Abstract

Computational simulations of natural selection for communicative success show that natural selection alone is capable of evolving optimal communication systems. Simulations of the interactions between natural selection and learning show that the biases of learners, when placed in the framework of iterated cultural transmission of communication systems, result in cultural selection of communication systems. This cultural selection is the determining factor in the development of communication systems in the simulated populations, with natural selection being relegated to a secondary role. This paper suggests that the role of cultural selection in the evolution of language should not be underestimated.

1 Introduction

Language is transmitted from generation to generation within a speech community. This transmission process involves at least some cultural transmission - language learners, under normal circumstances, learn the language in use in their speech community. The most influential linguistic theories of modern times (Chomsky, 1965; Chomsky, 1980; Chomsky, 1981; Chomsky, 1987) assume an element of genetic transmission in addition to this cultural transmission - under the Chomskyan hypothesis, a genetically-encoded Language Acquisition Device is transmitted from generation to generation, with the cultural transmission of language being parasitic on this innate language faculty.

The research outlined in this paper represents an attempt to understand the type of interactions which may occur between cultural transmission and genetic transmission during the repeated transmission of a communication system within a communicating population. In dealing with such complex, dynamical processes, the consequences of initial assumptions are not always clear. Computational modelling techniques allow the behaviour of complex systems to be monitored in detail over time, allowing the computational modeller to observe the consequences of initial assumptions in addition to forcing the modeller to make those assumptions explicit. For these reasons, the research outlined in this paper is informed by the use of computational simulation techniques.

A large body of computational modelling work has demonstrated that non-random genetic transmission (i.e. natural selection) of innate communication systems is capable of producing optimal, innate communication systems (e.g. Werner and Dyer (1991), Ackley and Littman (1994), MacLennan and Burghardt (1994), Levin (1995), Cangelosi and Parisi (1996), Oliphant (1996), Bullock (1997), de Bourcier and Wheeler (1997), Di Paolo (1997), Werner and Todd (1997), Noble (1998)). A growing body of computational modelling work suggests that iterated learning alone is capable of developing optimal, learned communication systems (e.g. Hutchins and Hazelhurst (1995), Steels and Vogt (1997), Batali (1998), Batali (in press), Livingstone and Fyfe (1999), Hurford (in press), Kirby (in pressb), Kirby (in pressa), Oliphant (in press)). A smaller body of work suggests that natural selection interacts in a positive manner with cultural transmission of communication systems to produce communication systems which are, in some sense, part innate and part learned (e.g. Batali (1994), Briscoe (1997), Kirby and Hurford (1997)). This paper explores a limited set of interactions between natural selection and cultural transmission in the evolution of simple communication systems and suggests that cultural transmission may be the determining factor in the development of communication systems in the simulated populations.

2 The communication system

An extremely simple model of communication is used in all simulations outlined in this paper, to focus inquiry on the interactions between natural selection and cultural transmission, rather than on the model of communication. Communication systems, also known as signalling systems, consist of a set of meaning-signal pairs. The canonical example of such a communication system is the Vervet monkey alarm call system (Cheney and Seyfarth, 1990). Such communication systems clearly lack syntax - signals are discrete tokens which may not be combined with other signals to form more complex signals. Optimal communication between two individuals using such a system requires that the individual wishing to communicate meaning M uses signal S, and the individual receiving signal S interprets it as meaning M.

Throughout this paper communication systems will be discussed in terms of the degree of *homonymy* they exhibit. Homonyms are "words having the same spelling and pronunciation but differing in meaning" (The Concise Oxford Dictionary). Homonymous signals in these communication systems therefore introduce ambiguity as there is no syntactic context from which to identify the intended meaning of the homonymous signal[1]. Communication systems will be termed:

- *Unhomonymous* if every meaning is associated with a distinct signal. The Vervet communication system is an example of an unhomonymous communication system.
- *Partially homonymous* if some, but not all, meanings are associated with identical signals.
- *Fully homonymous* if all meanings are associated with identical signals. In such systems, all signals will be homonyms of all other signals and therefore maximally ambiguous.

3 Communicative agents

Small feedforward neural networks are used to model communicative agents capable of communicating using the type of communication systems outlined above. Feedforward networks map an input pattern of activation through a set of weighted connections to produce an output pattern of activation. For the purposes of the simulations outlined in this paper, input patterns are considered to be meanings and output patterns are considered to be signals. The structure of the neural network used is shown in Figure 1.

Binary input meanings are used, and thresholding of the output of the neural network results in binary signals. Communication systems therefore consist of mappings from binary meanings to binary signals. The actual mapping from meanings to signals produced by a network is determined by the relative

[1] Homonymy is rife in natural languages, but utterances containing homonymous words are rarely ambiguous due to the context provided by the rest of the utterance. In English, for example, "bank" has several interpretations - it can refer to a river's edge or a financial institution or the act of tilting sideways while making a turn, to name but a few. However, the utterance "I paid money into the bank" will not normally be ambiguous due to the linguistic context the homonymous "bank" appears in

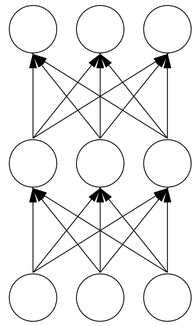

Figure 1: The neural network.

strengths of the connection weights in the network. Each agent's set of meaning-signal pairs is produced by propagating all meanings through their neural network to produce a signal for each meaning. This set of meaning-signal pairs is then used for production and reception. The production and reception processes are summarised below.

- *Production:* To communicate meaning M, use the signal S associated with M in the agent's meaning-signal inventory.
- *Reception:* On receiving S_i, retrieve the signal S_j from the agent's meaning-signal inventory most similar[2] to S_i and interpret S_i as communicating M associated with S_j in the agent's meaning-signal inventory. If there are several equally similar signals, select one randomly.

4 Simulating natural selection

The communication system used by a communicative agent of the type outlined above is determined by the weights of the connections in that agent's neural network. Previous computational modelling work suggests that, given a genetic encoding of these connection weights and selection pressure for communicative success, optimal, innate communication systems will emerge in the simulated population. A genetic algorithm (Holland, 1975) is used to simulate this process of natural selection.

4.1 Ingredients for natural selection

The simulation of natural selection outlined in this section has four key components:

1. A model of population turnover.

[2]Hamming distance is used as a measure of similarity between signals.

Figure 2: The generational model of population turnover.

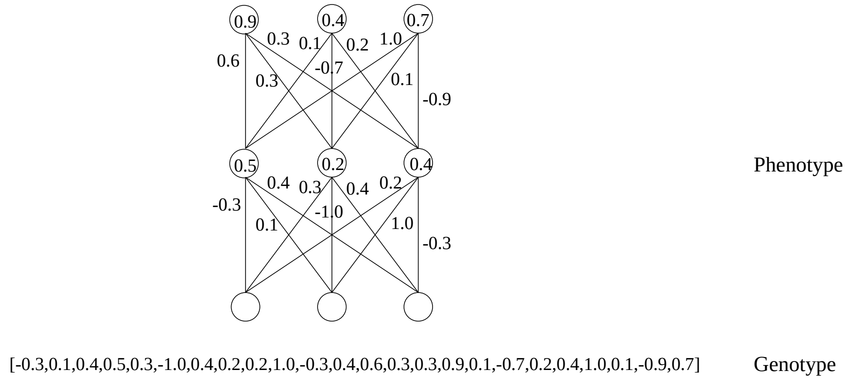

Figure 3: Part of the mapping from genotype (a string of real numbers) to phenotype (a neural network). Bias node connection weights are shown in the associated node.

2. A model of a communicative agent.
3. A genetic representation of connection weights.
4. A method of recombination of genes to produce new gene combinations.
5. Selective breeding based on an evaluation of communicative ability.

A generational model of population turnover is used - the simulation is organised into generations, each generation consisting of a complete population. Generation N breeds to produce generation $N+1$ and then dies off, leaving generation $N+1$ to breed to produce a new generation. This iterated process of breeding is illustrated in Figure 2.

The model of a communicative agent is the neural network model outlined above. The genetic representation of connection weights, the *genotype*, is a string of real numbers. Each position, or locus, on the genotype maps to a particular connection in the neural network, the *phenotype*. The value of the gene at the given locus becomes the weight of the associated connection in the neural network phenotype. This relationship between genes and neural network connection weights is depicted graphically in Figure 3.

Genetic transmission is facilitated by breeding - agents breed to produce new agents. The new agent's genes are a combination of the genes of the parent agents. Mutations also occur during the genetic transmission process - with a small, fixed probability single genes on the offspring genotype are incremented or decremented by a small amount.

- For each agent in the population:
  - Call that agent the speaker.
  - For each of the speaker's neighbours, do:
    - Call that neighbour the hearer
    - For each meaning, M_s, do:
      - Select the signal, S_s, that the speaker associates with M_s
      - Identify the meaning, M_h, that the hearer interprets S_s as conveying using the reception process summarised above.
      - Compare M_h with M_s and score the success of the communication. If M_h is identical to M_s score the communication as a success. Otherwise, the communication is a failure.
    - Return the hearer to the population.
  - Return the speaker to the population

Figure 4: The algorithm for evaluating communicative success of agents in the population.

These first four elements introduce the mechanisms of genetic transmission into the simulation. In order to model selection for communicative success, selective breeding based on an evaluation of communicative success is included in the model. The population of agents is organised on a toroidal line - each agent in the population has two neighbours on this line. Each agent is evaluated in terms of how successfully they communicate with their immediate neighbours, with agents who are successful communicators being more likely to breed than agents who are unsuccessful communicators. The evaluation algorithm is summarised in Figure 4. Although the population is spatially organised for the purpose of evaluation, the population is not spatially organised for breeding - breeding is not restricted to neighbouring agents, and the offspring of two agents is placed in a random position in the next generation. This prevents speciation in the population, a process which is outside this paper's focus of inquiry.

4.2 The results of natural selection

The average communicative success of members of 10 communicating populations is shown in Figure 5. The communication system types in use, averaged over these 10 populations, are shown in Figure 6. For all the simulations outlined in this paper, the agents had three meanings to communicate. 33% communicative success therefore corresponds to a population fully converged on a fully homonymous communication system, 66% success corresponds to a population converged on a partially homonymous system, and 100% success corresponds to a population converged on an optimal, unhomonymous communication system.

As can be seen from Figures 5 and 6, natural selection has proved to be

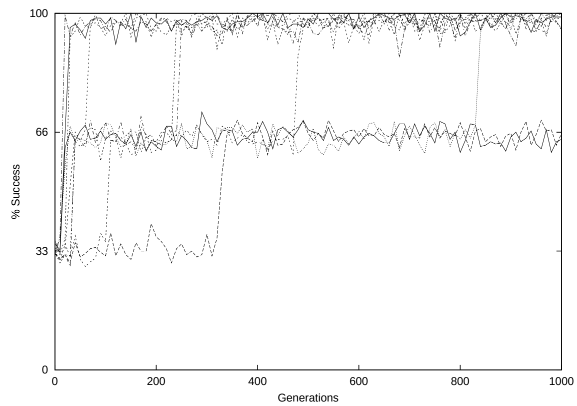

Figure 5: Average communicative success of members of 10 populations over time.

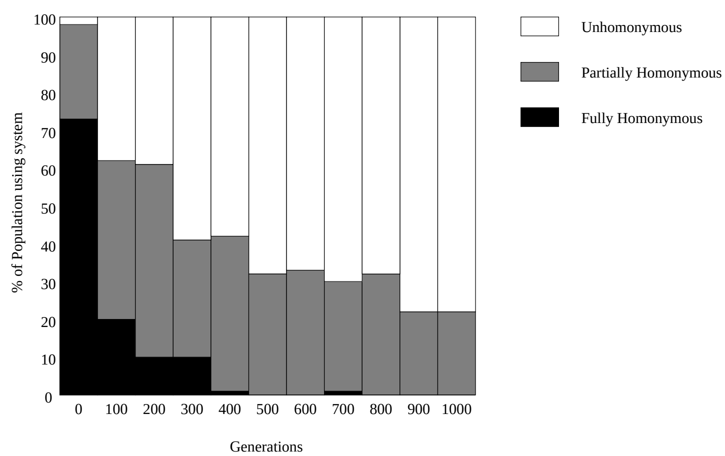

Figure 6: Communication system types in use over time, averaged over 10 populations.

effective at evolving optimal, innate communication systems in the majority of the simulated populations. This finding is in agreement with previous computational modelling work on the evolution of innate communication systems.

Natural selection is "the differential survival and reproduction of ... different genotype[s] within a population" (Jones, Martin, and Pilbeam, 1992, p467). In the simulations outlined in this paper, genotypes are differentiated on the basis of the communicative success of the phenotypes they encode - genes which are expressed as phenotypes which enjoy high communicative success are more likely to be retained in the population's gene-pool - successful communicators breed and pass on their genes. This gradual process of filtering of genes eventually results in the majority of the populations shown in Figures 5 and 6 consisting almost entirely of agents whose genotypes encode unhomonymous communication systems. Differential retention of genotypes is the driving mechanism of natural selection.

5 Adding cultural transmission

In the simulations outlined in the previous section, the only form of intergenerational transmission was genetic transmission, guided by selection for communicative success. Previous work on the interactions between genetic transmission and cultural transmission suggests that the addition of cultural transmission will result in positive interactions between the two forms of intergenerational transmission, such as the Baldwin Effect (Baldwin, 1896). In this section the addition of cultural transmission to the computational model is described.

5.1 Learning

Intergenerational cultural transmission is facilitated by learning - in the simulations outlined in this section, individual agents not only inherit their genes from the preceding generation, but also observe and learn from the communicative behaviour of the preceding generation. All the ingredients required for simulating natural selection, outlined in section 4.1, are included in the new simulation - a model of population turnover, adjusted to include cultural transmission (see Figure 7), a model of a communicative agent, a genetic representation of connection weights, a method of recombination of genes and selective breeding. In the previous section, an agent's communication system was entirely determined by the connection weights encoded in their inherited genes. In the simulations involving cultural transmission, these inherited connection weights merely form the starting point for learning a communication system.

The population at generation N and their offspring population, generation $N+1$, are spatially organised relative to one another - both populations are organised in a toroidal line, and certain positions in the generation N population are closer to certain positions in the $N+1$ population. An agent in generation $N+1$ learns from the three agents in generation N closest to them. This is illustrated graphically in Figure 8.

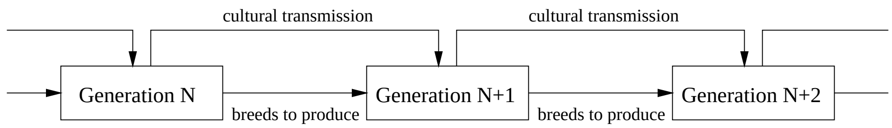

Figure 7: The model of population turnover and cultural transmission.

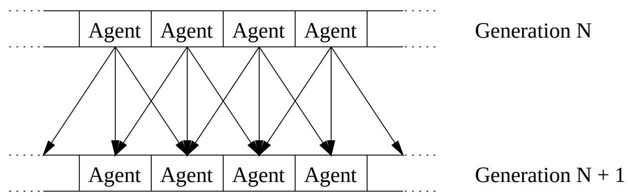

Figure 8: Spatial organisation of generation N and $N+1$. Arrows indicate cultural transmission of communication systems.

Each individual in the population at generation $N+1$ receives 25 exposures to the communication systems of the population at generation N. These exposures are randomly distributed among the three closest members of generation N. During each exposure, the set of meaning-signal pairs of the N generation agent are used to train the generation $N+1$ agent. The backpropagation algorithm, a neural network learning algorithm which makes adjustment to network connection weights, was used to implement this learning process, with the starting point for learning being the connection weights specified in the learning agent's genotype. The communication system of a generation $N+1$ agent is therefore determined by the interactions between genetic and cultural transmission of the communication systems of a proportion of the previous generation.

5.2 Cultural transmission and cultural stagnation

The average communicative success of members of 10 learning populations over time is shown in Figure 9. Figure 10 illustrates the proportion of communication systems of the three types in use over time, averaged over the 10 learning populations. It is clear from Figures 9 and 10 that the addition of cultural transmission has prevented the populations from developing optimal communication systems. In fact, the populations have converged on fully homonymous systems, the worst type of system.

The poor performance of the learning populations is due to the biases of the learners. There are two possible biases which may effect the successful transmission of communication systems - learning biases making certain systems harder to learn than others, and learning biases determining how the learners respond when confronted with conflicting communication systems to learn from.

Figure 11 summarises the success of networks with small, random initial

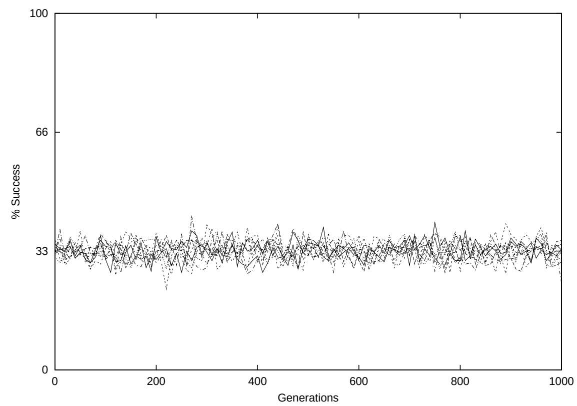

Figure 9: Average communicative success of members of 10 learning populations over time.

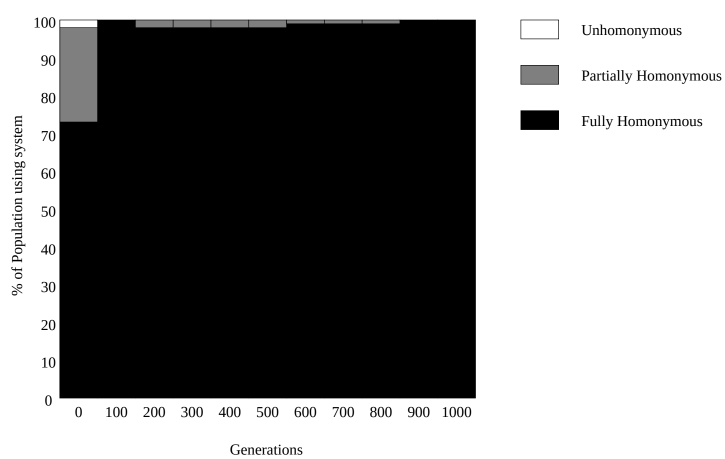

Figure 10: Communication system types in use over time, averaged over 10 learning populations.

System Type	% success
Unhomonymous	99.5
Partially Homonymous	100
Fully Homonymous	100

Figure 11: The success of networks at the learning task. 100 random communication systems of each type were each tested on 100 networks.

System Type	% population
Unhomonymous	1.8
Partially Homonymous	26.4
Fully Homonymous	71.8

Figure 12: The percentage of a population of agents with small, random weights using communication systems of the given type

weights in acquiring communication systems of the three types given 25 exposures to single systems of those types. This reflects the ability of the communicative agents to acquire a system of the given type when presented with unified communicative behaviour to learn from. As can be seen from Figure 11, the agents enjoy extremely high success at acquiring systems of all three types in these circumstances, with unhomonymous systems proving fractionally harder to acquire. This small learning bias is clearly not the determining factor in the evolving communicative behaviour of the learning populations.

Figure 12 summarises the percentage of networks with random initial connection weights using communication systems of the three types. This approximates the response of learning agents to training on conflicting communication systems - training on conflicting systems effectively randomises the connection weights in the agents' neural networks. As can be seen from Figure 12, under these circumstances the majority of agents are likely to converge on fully homonymous communication systems. This is the crucial learning bias which results in the convergence of the learning populations on fully homonymous communication systems.

As can be seen from Figure 10, there is a mix of communication systems in the populations at generation 0. Agents at generation 1 therefore observe and learn from a mix of communication systems. Under such circumstances, the agents are likely to acquire fully homonymous systems, due to the learning bias in favour of acquiring fully homonymous systems when presented with mixed communicative behaviour. The same will be true of learners in subsequent generations faced with mixed communicative behaviour. The learning biases of the agents filter out communication systems not conforming to those biases through repeated iterations of cultural transmission.

In Section 4.2, natural selection was characterised as "the differential survival and reproduction of ... different genotype[s] within a population". This differential reproduction of genotypes is a result of properties of the genotypes themselves - some genotypes have properties which make them more likely to be genetically transmitted through reproduction. In the simulations outlined in this paper, genotypes which encode neural networks which are successful communicators are more likely to be passed on through breeding. Repeated filtering of genotypes by selective breeding results in the evolving populations coming to be dominated by genotypes which encode optimally communicating neural networks - neural networks which conform to the breeding bias in favour of successful communicators.

Similarly, some communication systems are more likely than others to be successfully transmitted between generations. This selective transmission of communication systems is determined by the learning biases of the learning agents, with communication systems conforming to the learning biases being more likely to be successfully transmitted. This differential reproduction of different communication systems within the population can be termed *cultural selection*. In the learning populations outlined in this section, the learning biases of the agents result in cultural selection for fully homonymous communication systems and as a result these systems come to dominate the populations.

Why does natural selection not counteract the force of cultural selection and weed out poor communicators? Learning in the phenotype masks an individual's genetic makeup - no matter how good an agent's genes are, their effects are likely to be overtaken by learning, which almost fully determines an agent's communicative behaviour. Shielding (Ackley and Littman, 1991) prevents natural selection from identifying good gene combinations and weeding out bad gene combinations. Even if an agent were to inherit genes which were so good they could not possibly learn a fully homonymous system, they would be no better off - an agent must communicate with its neighbours, and if those neighbours use a fully homonymous system then using a better communication system to communicate with them yields no benefit. Cultural transmission leads to cultural stagnation in the simulated populations - cultural selection favours fully homonymous communication systems and natural selection is powerless to counteract this.

5.3 Cultural transmission and collapse

The addition of cultural transmission and the resultant cultural selection of communication systems not only prevents the development of optimal communication systems in the simulated populations - it also prevents the maintenance of such a system. Figure 13 shows the average communicative success of a population of agents who initially have a shared, optimal, innate communication system - all the agents in the population at generation 0 have a set of genes which encode an unhomonymous communication system. As in the other simulations outlined in this section, every generation after the first attempts to learn a communication system based on the communicative behaviour of the preceding generation. Figure 13 clearly shows the population collapsing from using an unhomonymous communication system to using a fully homonymous

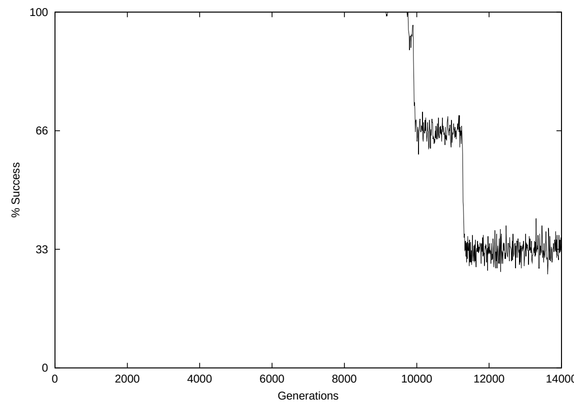

Figure 13: Average communicative success of members of a learning population over time.

communication system within 12000 generations.

As mentioned above, learning almost completely masks an agent's genetic makeup. However, there are certain gene combinations which make learning a given communication system impossible - an agent's genes constitute the starting point for learning, and the backpropagation learning algorithm is sensitive to starting weights to a certain extent. Shielding of genes from natural selection results in genetic drift - learning prevents natural selection from evaluating genes, so genetic transmission is essentially random. As a result of this genetic drift, an agent in the population shown in Figure 13 will eventually be born whose genes are so bad that it cannot acquire the unhomonymous communication system in use by the rest of the population, and will acquire a partially or fully homonymous communication system instead. This individual will be unlikely to breed, due to low communicative success with its neighbours, but its communication system will be culturally transmitted to the next generation.

The agents in the next generation observing this suboptimal communication system will also observe a sample of the optimal communication systems present in the previous generation. As discussed above, under such circumstances these agents may acquire suboptimal communication systems. They will not breed, but their communication systems will be passed on culturally to subsequent generations. Cultural selection favouring increasing levels of homonymy allows suboptimal systems to spread through the population like a virus, resulting in the collapse in communicative success shown in Figure 13. This spread of suboptimal systems through the population is illustrated in Figure 14.

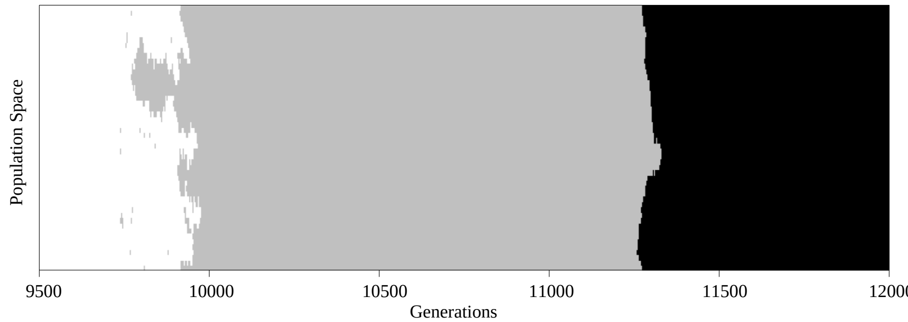

Figure 14: The population collapses from using an unhomonymous system (white) to a partially homonymous system (grey), and then to a fully homonymous system (black). over 10 populations.

6 Conclusions

This paper outlines a computational model of the emergence of communication in a simulated population of communicative agents. As previous work suggests, natural selection alone is capable of developing optimal, innate communication systems in such populations. However, the addition of cultural transmission of communication systems prevents the emergence or maintenance of such systems. This is due to the shielding of genetic information by learning and the process of cultural selection during cultural transmission. Cultural selection is a result of the biases of the learners involved in the cultural transmission process - communication systems which conform to these biases are more likely to be successfully transmitted than communication systems which do not.

The simulations outlined in this paper give a clear indication of the fact that individual learner biases can have profound effects on the development of the communication system in use by a population when placed in the context of iterated cultural transmission. This suggests that the importance of cultural transmission and cultural selection should not be underestimated by researchers interested in the origins and evolution of language. This is not to say, however, that natural selection has no role to play in the evolution of language. In the simulations outlined in this paper, natural selection was restricted to tinkering with the starting point for learning, while the learning algorithm and the concomitant learning biases were externally imposed. Allowing natural selection to develop learning algorithms, and therefore exercise a degree of control over the nature of cultural selection on cultural transmission, would no doubt result in complex and interesting interactions between natural selection and cultural selection. Computational modelling will prove an invaluable tool in investigating such complex dynamic processes and developing a coherent theory of the evolution of language.

References

Ackley, D. and M. Littman (1991). Interactions between learning and evolution. In C. Langton, C. Taylor, J. Farmer, and S. Rasmussen (Eds.), *Artificial Life 2*, pp. 487–509. Redwood City, CA: Addison-Wesley.

Ackley, D. and M. Littman (1994). Altruism in the evolution of communication. In R. Brooks and P. Maes (Eds.), *Artificial Life 4: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pp. 40–48. Redwood City, CA: Addison-Wesley.

Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist 30*, 441–451.

Batali, J. (1994). Innate biases and critical periods: Combining evolution and learning in the acquisition of syntax. In R. Brooks and P. Maes (Eds.), *Artificial Life 4: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pp. 160–171. Redwood City, CA: Addison-Wesley.

Batali, J. (1998). Computational simulations of the emergence of grammar. Cambridge: Cambridge University Press.

Batali, J. (in press). The negotiation and acquisition of recursive grammars as a result of competition among exemplars. In E. Briscoe (Ed.), *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press.

Briscoe, E. (1997). Language acquisition: the bioprogram hypothesis and the baldwin effect. MS, Computer Laboratory, University of Cambridge.

Bullock, S. (1997). An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models. In P. Husbands and I. Harvey (Eds.), *Fourth European Conference on Artificial Life*, pp. 454–463. Cambridge, MA: MIT Press.

Cangelosi, A. and D. Parisi (1996). The emergence of a 'language' in an evolving population of neural networks. Technical Report NSAL-96-004, Institute of Psychology, National Research Council, Rome.

Cheney, D. and R. Seyfarth (1990). *How Monkeys See the World: Inside the Mind of Another Species*. Chicago, IL: University of Chicago Press.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Chomsky, N. (1980). *Rules and Representations*. London: Basil Blackwell.

Chomsky, N. (1981). *Government and Binding*. Dordrecht: Foris.

Chomsky, N. (1987). *Knowledge of Language: Its Nature, Origin and Use*. Dordrecht: Foris.

de Bourcier, P. and M. Wheeler (1997). The truth is out there: The evolution of reliability in aggresive communication systems. In P. Husbands and

I. Harvey (Eds.), *Fourth European Conference on Artificial Life*, pp. 444–453. Cambridge, MA: MIT Press.

Di Paolo, E. (1997). An investigation into the evolution of communication. *Adaptive Behaviour 6*, 285–324.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems.* Cambridge,MA: MIT Press.

Hurford, J. R. (in press). Expression/induction models of language evolution: Dimensions and issues. In E. Briscoe (Ed.), *Linguistic Evolution through Language Acquisition: Formal and Computational Models.* Cambridge: Cambridge University Press.

Hutchins, E. and B. Hazelhurst (1995). How to invent a lexicon: The development of shared symbols in interaction. In N. Gilbert and R. Conte (Eds.), *Artificial Societies: the computer simulation of social life.* London: UCL Press.

Jones, S., M. Martin, and D. Pilbeam (Eds.) (1992). *The Cambridge Encyclopedia of Human Evolution.* Cambridge: Cambridge University Press.

Kirby, S. (in pressa). Learning, bottlenecks and the evolution of recursive syntax. In E. Briscoe (Ed.), *Linguistic Evolution through Language Acquisition: Formal and Computational Models.* Cambridge: Cambridge University Press.

Kirby, S. (in pressb). Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. In C. Knight, M. Studdert-Kennedy, and J. R. Hurford (Eds.), *The Evolutionary Emergence of Language.* Cambridge: Cambridge University Press.

Kirby, S. and J. R. Hurford (1997). Learning, culture and evolution in the origin of linguistic constraints. In P. Husbands and I. Harvey (Eds.), *Fourth European Conference on Artificial Life*, pp. 493–502. Cambridge, MA: MIT Press.

Levin, M. (1995). The evolution of understanding: a genetic algorithm model of the evolution of communication. *BioSystems 36*, 167–178.

Livingstone, D. and C. Fyfe (1999). Modelling the evolution of linguistic diversity. In D. Floreano, J. Nicoud, and F. Mondada (Eds.), *Advances in Artificial Life: Fifth European Conference on Artificial Life*, pp. 704–708. Berlin: Springer.

MacLennan, B. and G. Burghardt (1994). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behaviour 2*, 161–187.

Noble, J. (1998). Evolved signals: Expensive hype vs. conspiratorial whispers. In C. Adami, R. Belew, H. Kitano, and C. Taylor (Eds.), *Artificial Life 6: Proceedings of the Sixth International Conference on Artificial Life.* Cambridge, MA: MIT Press.

Oliphant, M. (1996). The dilemma of saussurean communication. *BioSystems 37*, 31–38.

Oliphant, M. (in press). The learning barrier: Moving from innate to learned systems of communication. To appear in *Adaptive Behaviour.*

Steels, L. and P. Vogt (1997). Grounding adaptive language games in robotic agents. In P. Husbands and I. Harvey (Eds.), *Fourth European Conference on Artificial Life*, pp. 474–482. Cambridge, MA: MIT Press.

Werner, G. and M. Dyer (1991). Evolution of communication in artificial organisms. In C. Langton, C. Taylor, J. Farmer, and S. Rasmussen (Eds.), *Artificial Life 2*, pp. 659–687. Redwood City, CA: Addison-Wesley.

Werner, G. and P. Todd (1997). Too many love songs: Sexual selection and the evolution of communication. In P. Husbands and I. Harvey (Eds.), *Fourth European Conference on Artificial Life*, pp. 434–443. Cambridge, MA: MIT Press.