

Evolution and Learning of Language: Insights Drawn from Modeling

M. Sean McLennan
P717 Evolution and Learning – Dr. William Timberlake
Indiana University, Bloomington

1 Introduction

How can we characterize the dual contributions of biology and environment to language? This is one of the most hotly contested topics in linguistics, psychology, and related disciplines. Even the staunchest adherents to both ends of the “nature / nurture” spectrum must acknowledge that there are both components of genetics and learning to language: no one would claim that specific languages are encoded in genes, nor would anyone propose that humans learn to have brains. The question becomes where in that spectrum we draw the line.

Strong nativists (Pinker & Bloom (1990) for example) hold that even aspects of grammar are genetically determined and thus can be selected for via natural selection. This accounts for the kind of cross-linguistic regularity that characterizes “Universal Grammar” – the linguistic knowledge that all humans possess. On the other end, strong developmentalists would maintain that the creation of neurons is as far as biology contributes; before the contribution of environment, the brain is just a ball of “connectoplasm”¹. Any regularity in language is the result of regularity in the environment and other physiological regularity (for example, we all have eyes and ears that respond to stimuli in highly structured and similar ways). Thelen & Smith (1993) is representative of the approach taken by researchers on this end of the spectrum.

¹This vivid term, “connectoplasm” is intended to allude to “cytoplasm”, and means a unstructured, homogeneous mass of neurons.

Nativists refer to reports like Gopnick (1990) which describes the KE family who show a language deficit that appears to be caused by a recessive gene. Developmentalists point to the profound cortical plasticity of the brain and its ability to self-organize; for example, within the visual system which is relatively well understood (Gazzaniga et al. 1998, pg 123-134, 484-493). As with most arguments, the truth likely lies somewhere in the middle; from the preface of Elman et al. (1996: xi-xii):

“The obvious conclusion is that the real answer to the question, *Where does knowledge come from*, is that it comes from the *interaction* between nature and nurture, or what has been called ‘epigenesis.’ Genetic constraints interact with the internal and external environmental influences, and they jointly give rise to the phenotype.”

However, the authors go on to rightly point out: “Unfortunately, as compelling and sensible as this claim seems, it is less a conclusion than a starting point. The problem does not go away, it is simply rephrased.” (Elman et al. 1996: xii). No line has been drawn.

Particularly in studying the *evolution* of language and other cognitive capabilities, researchers must address the question of whether their pursuit is viable in the face of a profound lack of information. We have no way of obtaining direct information about the cognition of our ancestors. Clearly, Lewontin (1998: 128-129) has little respect for the study:

“Wherever cognition came from, one would like to make arguments about the forces of natural selection that established it. On the face of it, generalized problem solving and linguistic competence might seem obviously to give a selective advantage to their possessors. ... stories about how the species as a whole would be benefited by speech are not to the point. ... The problem is that we will never know [the factors involved] and never will. There is no end to plausible storytelling.”

Lewontin’s attack on the study of evolution of cognition is so vehement as to border on the belligerent. He continues to conclude (pg 130):

“Finally, I must say that the best lesson our readers can learn is to give up the childish notion that everything that is interesting about nature can be understood. History, and evolution is a form of history, simply does not leave sufficient traces, especially when it is the forces that

are at issue. Form and even behavior may leave fossil remains, but forces like natural selection do not. It might be interesting to know how cognition (whatever that is) arose and spread and changed, but we cannot know. Tough luck.”

Despite Lewontin’s overly emotional, and perhaps condescending attitude, his point is valid. Without evidence that is more conclusive than speculative, studying the evolution of cognition is no more than telling “just-so” stories.

Botha (1997a, 1997b, 1998a, 1998b., 2000), in a series of articles that discusses Neo-Darwinian accounts of language, particularly focusing on Pinker & Bloom (1990), systematically criticizes both Pinker & Bloom (1990) but also their detractors. Botha’s point is not to take a stance on either side of the issue, but to point out some of the fundamental flaws that hamper the proper exposition of the arguments. Perhaps the most damning comes from Botha (1997a) in which he systematically shows that no one involved in the argument is discussing the same ontological entity. Even within Pinker and Bloom (1990), the evolving linguistic entity under scrutiny varies wildly (Botha, 1997a: 253).

With such criticisms being brought to bear against the entire domain of study, the question moves away from “where do we draw the line between nature and nurture?” to “can examining the evolution of language lead to a better understanding of language?” There is one area of language evolution research that seems to adequately acknowledge evidential restrictions, refrain from “just-so” storytelling, and provide profound insights into the nature of the interaction between language evolution and language learning. That area is modeling, the primary focus of this discussion.

It is worth quoting the introduction to Batali (1998) at length for it addresses (at least indirectly) many of the issues brought forth above:

“The ability to communicate by exploiting a system of structural regularities ... represents an invaluable achievement of a species for which coordinated social activity is vital

to survival, and for which the accurate exchange of information often provides adaptive benefit. Given this benefit, it is **tempting** to explain the achievement as the result of natural selection.

“However it is important to distinguish between the evolution of language itself – in particular the emergence, modification, and enrichment of the grammatical resources in human languages – **and the biological evolution of articulate hominids.** Clearly they are related: Lacking appropriate anatomical and neurological endowment, an animal will be unable to produce or perceive complex signals, and without sufficiently powerful cognitive abilities, it cannot entertain meanings worth communicating in the first place. But the adaptive benefits of such traits are not specific to communication, and it is not clear how communication alone could provide sufficient selection pressure for their development.

“In this paper I explore the idea that some of the grammatical regularities manifest in human language could emerge as a result of non-genetic cultural processes among a population of animals with the cognitive capacities required for communication, but who do not initially share a coordinated communication system. **Whether or not this is what really happened in our species is unknown, but the possibility seems worth investigating, to better understand both its plausibility and its limitation.**

“Speculation about the origins and early development of human language must perforce originate in the intuitions based on experience with its modern versions. Though these intuitions can be tested methodologically, with the resources of linguistics, psychology, anthropology, neuroscience, literary studies, and other disciplines, we are for the most part limited to the one class of exemplars. There is obviously no way to go back and observe what happened since the Pliocene epoch, and the differences between the regularities exhibited by the grammars of human languages and those of other animal communication systems seem profound.

“Mathematical and computational models provide a way to explore alternative accounts of the emergence of systems of communication. If the consequences of a model are consistent with expectations based on intuitions or speculation, they might obtain some **small** measure of support. But more interestingly (and, as it happens, more often), the consequences of a model may deviate from expectations. In working out the reasons for the differences, one can potentially develop a richer set of intuitions. **Models are thus valuable to the degree that they explicitly illustrate the consequences of the set of assumptions they embody. This may be even more important than whether those assumptions are correct.**”

[Emphasis added – MSM]

The attitude extant in this passage is one that is common among modelers of all varieties, and one that seems poorly understood by other disciplines. That is, that models are to be used as a tool for exploring assumptions, consequences, and the mapping of input to results. Indeed, it is a fundamental caveat of modeling that regardless of how well the model matches the data, never make

the leap from “model as description” to “model as explanation”.²

What follows is a review of three models of the evolution and learning of language and what they have brought to our understanding of the roles of biology and environment in language acquisition: Werner and Dyer’s (1991) model of the evolution of signaling systems, Batali’s (1998) model of the emergence of combinatorial grammar, and Kirby and Hurford’s (1997) model of the emergence of a “Language Acquisition Device”. I hope to make it clear that although discussions of the evolution of language are rightly suspect, it is not, as Lewontin would suggest, a “childish” pursuit.

2 Evolution of Communication in Artificial Organisms

The goal of Gregory Werner and Michael Dyer’s (1991) model was to explore the evolution of signaling systems within populations of organisms that had (via innate architecture) the ability (and impetus) to produce, perceive, and respond to signals.

2.1 Architecture

The model was implemented within an Artificial-Life (A-Life) framework: a population of 1600 artificial agents of two distinct types (labeled “male” and “female” for ease of reference) were introduced to a 200x200 toroidal grid / world. Each agent has a distinct genome that is expressed

²This is a point it would behoove Linguists in general to take note of. In a very real sense, modern linguistic theories such as generative grammar and universal grammar are models of language yet it is overwhelmingly assumed that they have explanatory power and direct physical correlates in biology and neurology. Cook & Newson (1996: 32) state: “The principles of UG should be relatable to physical aspects of the brain; the brain sciences need to search for physical counterparts for the mental abstractions of UG” in discussing Chomsky’s (1986: 39) comment, “the abstract study of states of the language faculty should formulate properties to be explained by the theory of the brain”.

It is clear that this is a vital assumption in Pinker and Bloom’s (1990) argument and could possibly account for a great deal of the controversy surrounding it.

to create a recurrent neural network that controls its actions during its “lifetime”. It is important to emphasize that no learning occurs during the lifetime of the organism – the networks are static and fully specified by the genome.

Two general principles were considered in the designing the model: 1) there should be no direct pressure on the individuals to communicate; and 2) the populations should be presented with natural tasks that directly and predictably affect reproductive success, such as finding food, protecting young, and attracting mates. Thus, the task implemented in this model is “mate-finding”; a problem that can be benefitted by communication, but does not a priori require communication.

Females: Females in the model are stationary but have a limited “vision” of 2 squares in any direction – she is at the center of a 5x5 “visual field”. While a male is in that visual field, she emits a signal at each time step. The input to the females’ network is the state of her visual field, and the output is a 3 bit signal (with $2^3 = 8$ possibilities).

Males: Males are mobile but blind. They have no method of directly determining the location of females, however; they can receive the signals emitted by females when they are within a female’s visual range, and they react to those signals. The signals the females emit and the males’ reactions to those signals are initially random – i.e. at the beginning of the simulation, the signals are meaningless and are just as likely to result in a negative reaction from the male as a positive one. The input to the males’ network is the incoming signal (or lack of one) and the output is one of 4 possible actions: move forward, turn left, turn right, or stand still.

The differences in abilities and behavior of males and females was to prevent more simplistic solutions from being exploited. For example, if males were provided with the ability to localize the signal source they may evolve to simply “home in” on the signal instead of interpreting it. Similarly,

if females were permitted to move, they might simply learn to move towards the males that they can see. Within this framework, the most efficient reproductive method is for the female to direct the male to herself.

2.2 Procedure

An initial random population is created and placed (again randomly) in the grid. At each time step the state of the females' visual fields are entered as input to their networks and activation is cycled to produce an output signal. The signals are copied to males within the range of each female and cycled to produce a reaction. The positions / orientations of each male are updated and the time step is complete.

Reproduction occurs when a male finds a female, that is, when a male lands on the same square as the female. The genomes of each are combined using crossover and mutation to create two offspring which replace two other randomly selected agents in the world. Thus the population remains constant and organisms reproduce as soon as they prove their fitness instead of fitness being compared at fixed intervals.³ The offspring and parents are randomly placed back in the world to prevent them from continually reproducing immediately.

2.3 Results

Over the course of multiple runs definite stages of development emerged as the agents exploited increasingly effective behaviors:

³This is a large departure from traditional genetic algorithms in with generations are discrete and the population's fitness it simultaneously judged.

1. Initially, all males wandered randomly, and females signaled oblivious to their surroundings.
2. Males that stand still became extinct since they never find females to reproduce with.
3. Males that go straight take over the population. This is a good strategy since statistically, there will be three or four females in any given row or column. It also avoids recovering previously explored areas.
4. Males appear that turn when in the same row or column as a female. This is the stage in which a co-evolved signaling convention begins to emerge.
5. Females evolve to use their existing signals in more and more situations and males learn to turn specifically *towards* the females when in the same row or column. Otherwise, they move forward by default. This is the most efficient strategy for females and males to find each other.

Of course, from run to run, the actual values of the signals are completely arbitrary and variable. Indeed, since there are more possible signals than are necessary to develop an effective strategy, there are many possible mappings from signal to reaction. Werner and Dyer go onto explore the development of different signaling “dialects”, an exposition of which is beyond the scope of this paper. Interested readers are referred to their original publication.

It is worth noting that across all runs, a signaling system never failed to emerge.

2.4 Contribution

At no point do Werner and Dyer make any claims that their model is an accurate simulation of any living organism nor do they claim that it is a simulation of language. In fact they quite clearly state

the opposite. So, then, what have Werner and Dyer demonstrated? They have shown that given agents with certain very general characteristics – the ability to produce signals and the ability to perceive signals – and some bit of information that is of mutual benefit to share, a communication system is not only likely to emerge, but guaranteed to emerge. It is not claimed that this is the only situation that will result in communication, nor that it is the way communication emerged in reality. But given the pervasiveness of communication systems in nature, the model supports our intuitions that communication must be something that can robustly and regularly develop. Additionally, even though animal communication and simple signaling systems are clearly not human-like language, it is widely accepted that they are related – that some how, progress was made from a similar pre-linguistic communication system to human language (Donald, 1993; Deacon, 1997, for example). To gain insights into the precursor to language, of course, is to gain insights into language itself.

In and of themselves, Werner and Dyer’s simulations do not mean very much. But in the context of a larger study of evolution of cognition, it provides an important step in understanding the general forces that act in simple situations and result in complex emergent phenomena. The obvious next step is to extend the model, exploring different architectures and different underlying assumption to see if more general principles and constraints can be extracted in order to mold the criteria by which we judge our language evolution “storytelling”.


3 Computational Simulations of the Emergence of Grammar

Batali (1998), the introduction of which was quoted so extensively above, developed a simulation that incorporated learning only, exploring the likelihood of language-like grammatical structure spontaneously appearing in genetically naive agents. His agents are simple neural networks that

produce sequences of tokens; initially there is no coordinated communication, but over time and through “cultural transmission” a negotiated communication system with some rather striking regularities emerges.

3.1 Agents

Batali’s agents are simple recurrent neural networks with a basic structure shown in fig. 1. The “character input” is a series of nodes that represent a finite set of letters, and the “meaning output” is a vector of 10



real numbers that are interpreted as semantic content. **Figure 1 Agent RNN architecture (Batali, 1998)**

The recurrent nature of the network (i.e. output is mapped back to input as “context”) allows the network to learn / produce sequences. Thus “utterances” can be linearly ordered strings, not just individual letters or simultaneously input combinations of letters.

Transmission of utterances occurs between two agents: one the speaker and one the listener. In each communicative episode, there is a “target meaning” (1s and 0s only) that is accessible to both agents; it can be thought of as the context in which the communication occurs – the shared experience that the message describes.

3.2 Sending & Receiving

The goal of the speaker is to produce a string of characters that will reproduce the target meaning as close as possible. Importantly, the utterance produced, is *also* interpreted by the *speaker* as a sent message. The ability to self-monitor utterances is integral to the simulation. The speaker’s sending

algorithm can be summarized as:

1. Initialize all network inputs to zero.
2. Choose a character to transmit:
 - a) for each character determine how many output values would be correct (correct = rounded activation – binary function)
 - b) transmit the most successful character
 - c) if there is a tie for most number of correct output values, choose character that produces the least error (error = root mean square of the difference between output and target – continuous function)
3. Send character to hearer.
4. The character is processed by hearer and speaker according to the receiving algorithm described below
5. After processing:
 - a) if all values in the sender's output are correct, stop.
 - b) if not correct, but sequence length exceeds the cutoff, stop.
 - c) otherwise repeat from 2.

It is important to note that the speakers are not trained to produce particular outputs. Learning occurs only in the interpretation process, which, since it occurs within the same network, also affects the sending behavior.

Receiving proceeds according to the following algorithm:

1. Initialize all network inputs to zero.
2. Process each character:
 - a) set appropriate input node to 1.0 and all others to 0.0
 - b) spread activation forward
 - c) copy hidden node values to context nodes
3. Upon completion, copy output nodes to meaning vector.

3.3 Negotiation

For communication to be successful, the members of a population must use an agreed upon set of token-to-meaning mappings. The negotiation of these mappings occurs in the simulation in rounds which proceed in the following steps:

1. Randomly select an agent from the population to be the hearer / learner

2. Repeat 10 times:
 - a) Randomly select an agent from the population to be the speaker / teacher
 - b) Train the hearer's network to interpret the speaker's utterances, each presented once in random order, using the backpropagation algorithm
 - c) Return speaker to the population
3. Return the learner to the population

A set of 100 possible meanings was used as the training set.

3.4 Results

Four quantitative criteria were used in evaluating the coordination of the population: correctness, error, distinctness, and length. Correctness and error (as described above) are comparisons of network output to target meaning; an output is "correct" if it matches the target when rounded, error is the root mean square of the difference between the output and the target. Distinctness is the




Figure 2: Record of a simulation run of a population of 30 agents. Plotted each round are (co) correctness, (er) error, (di) distinctness, and (le) length.

	<i>one</i>	<i>they</i>	<i>you</i>	<i>yall</i>	<i>yup</i>	<i>me</i>	<i>we</i>	<i>mip</i>	<i>yumi</i>	<i>all</i>
<i>tired</i>	cda	cdab	cdc	cdcb	cdba	cd	cdd	cddb	cdcd	cdb
<i>scared</i>	caa	caab	cac	cacb	caba	ca	cad	cadb	cacd	cab
<i>sick</i>	daa	daab	dac	dacb	daba	da	dad	dadb	dacd	dab
<i>happy</i>	baa	baab	bca	bcab	baac	ba	badc	bab	bac	babc
<i>sad</i>	aba	abab	ac	acb	abac	a	abdc	abb	abc	abbc
<i>excited</i>	cba	cbab	cca	ccab	cbca	c	ccdc	cb	ccb	cbc
<i>angry</i>	bb	bbb	bc	bcb	bbc	b	bddc	bdb	bdc	bdbc
<i>silly</i>	aa	aaab	aca	acab	adba	add	addc	adad	adc	adbc
<i>thirsty</i>	dbaa	dbab	dca	dcba	dbca	dda	ddac	dbad	dcad	dbacd
<i>hungry</i>	dbb	dbbd	dc	dcb	dbc	dd	ddc	dbd	dcd	dbcd

Figure 3: Sequences negotiated by a population for the range of possible meanings.

average fraction of utterances sent for exactly on meaning; i.e. a measure of the average number of utterances used to convey a single meaning (since multiple-to-one mappings are possible). Length is simply the average total length of utterance. A graph of the results of a run of the simulation is reproduced in fig. 2 (Batali, 1998). Over the course of over 15,000 rounds, an efficient communication system is negotiated which accurately conveys meaning while simultaneously minimizing the length of utterances and maximizing one-to-one mappings of form to meaning.

Moreover, if the utterances themselves are examined, a great deal of systematicity can be observed. For example, the paradigm in fig.3 was typical⁴. Although certainly not a perfectly coherent system, there are obvious regularities that border on the use of combinatorial rules. Subjectively, one is very strongly reminded of morphological paradigms in natural language.

Sensibly, Batali is quick to caution against such an analysis – the simulation clearly does not exhibit language. Instead of examining the sequences in terms of interpretation by the agents, he suggests that regularities emerge because the meanings share “trajectories through the vector space

⁴Batali imposed a completely arbitrary Subject-Predicate interpretation using English words on the meaning values. This was to facilitate examining the data since it is simpler than reading binary strings.

of the network activation values” – i.e. there are shared values in the meanings, and since the meanings are encoded in network weights, it is natural that to the extent that those values are shared, the form of the utterances that reproduce those values via the weights should also be shared.

The next question to be asked is naturally, “can the agents generalize?” Can they produce and interpret novel utterances they haven’t been trained on? In order to answer this question, Batali performed a simulation in which the agents were only trained on 90 of the 100 meanings and tested on the remaining 10. They performed reasonably accurately, and in a way that was generally coherent with the negotiated paradigm. Interestingly, the average length of novel utterances was slightly longer, which seems intuitively natural.

3.5 Contribution

As with Werner and Dyer’s model, we cannot make an direct conclusions about the evolution or acquisition of language, nor should we try. The model simply provides a framework in which to test assumptions and possibilities. This case is particularly interesting because it suggests that the some of the claims made by Pinker & Bloom (1990) are not necessarily correct. That it is possible that more linguistic structure than they believe can emerge as a direct consequence of learning and cultural transmission.

However, this is not to say that the results of this simulation are not compatible with Pinker & Bloom’s stance on the role of innateness in language, or even more generally the stance of Evolutionary Psychology. There are some very explicit assumptions made in Batali’s model. That agents have access to their own utterances and can self-monitor; that the speaker and hearer can share the same meaningful experience and that it is encoded in a similar way; that the architecture that

instantiates the learning is similar (if not identical) across individuals; that the speakers and hearers have an impetus to communicate and motivation to learn. All of these things are still very strong candidates for genetically encoded traits. The model only provides us with guidance in making educated guesses about the dynamic that is clearly prevalent between evolution and learning, and improve our “storytelling”.

4 Learning, Culture, and Evolution in the Origin of Linguistic Constraints

Modern linguistic theory more or less accepts that there is an innate language acquisition device (LAD) that all humans are born with (Chomsky, 1986). The details of what is universal (specified by the LAD) and what is not universal (not specified by the LAD) is the primary focus of the field of Linguistics. From an evolutionary perspective, this immediately begs the question, “Where did this LAD and its constraints come from?” This question is the focus of much of the work of Simon Kirby and James Hurford at the Language Evolution Center at the University of Edinburgh.

It seems clear that at least some “universals” of language serve a functional purpose: for example there is a syntactic constraint called “subjacency” that constrains movement dependent on the hierarchical structure of a sentence. It tends to eliminate sentences where there is a long gap between the moved element and its co-indexed original position – i.e. it tends to eliminate sentences that make large demands on memory. The most common evolutionary explanation for how this functional constraint became instantiated in the LAD is that the pressure to keep gaps at a minimum led to the natural selection of LADs that had a method of rejecting sentences with large gaps. Thus the constraint is a direct consequent of or innate LAD. Kirby and Hurford label this point of view “phylogenetic functionalism.”

Kirby and Hurford propose another explanation, which they term “glossogenetic functionalism”. It differs in that such constraints are not assumed to arise directly from the LAD. Universals instead, appear over both historical and cultural time-scales through acquisition and use – through learning and cultural transmission. It relies on the principle that the language of a learner does not necessarily converge on exactly the same system as the teacher.

4.1 The Model

Both the phylogenetic and glossogenetic explanations were explored by Kirby and Hurford (1997) within a computational simulation. The representation of “grammar” was an 8-bit string of 1s and 0s leading to $2^8 = 256$ possible “languages”. The representation of the LAD was similar – an 8-bit string – but with 3 possible alleles: 0, 1, and ? ($3^8 = 6561$ possible LADs). The ? represents an aspect of grammar that is not innate, but that can be set over the course of a lifetime depending on input. Consequently, the least constrained LAD – one with 8 ?s – could potentially learn all 256 possible languages, and the most constrained – with no ?s – exhibits no learning and is a completely innate instantiation of only 1 language.

In keeping with current linguistic theory about the setting of parameters in language, utterances are considered “triggers” that provide information about the appropriate grammatical settings. Utterances in the model are also 8-bit strings with three possible settings: 1, 0, or *. A 1 or a 0 means that the utterance carries information that can potentially act as a trigger for the grammatical relationship represented in that position. An * means that the utterance does not provide information about the target grammar to be learned. In this particular simulation, only utterances that provide evidence for one position in their grammar are produced (i.e. a 1 or 0 + seven *s). Of course,

utterances are consistent with the grammar of the speaker. Whether a particular trigger causes a parameter value to be set by a learner is a probabilistic function referred to as the trigger learning algorithm.

4.2 Critical Period and Selection

An important aspect of this model is the implementation of a “critical period” that is analogous to the critical period of language acquisition found in human children. Within the critical period, grammatical learning can occur, but not after the period has passed. Since it is possible that within that period, a sufficient number of triggers may not be observed by the learner, it is possible that grammars may not be completely / accurately transmitted even when it is within the range of possibilities specified by the LAD. This provides for variation in cultural transmission in the model – what Kirby and Hurford term “linguistic selection”.

Since this is an evolutionary model, a measure of communicative fitness needs to be applied to the agents of the population in order to model natural selection. This measure was taken during communicative episodes between adults after the critical period. Fitness was assessed based on both transmission ability and reception ability. Each agent is involved in a specified number of communication episodes, 50% as “hearer” and 50% as “speaker”. Each utterance produced is measured on whether it is analyzable by the hearer and how parsable it is. Analyzability is a binary function asking “is the utterance consistent with the hearer’s grammar”. Parsability is a non-binary function that, in this simulation, arbitrarily prefers 1s in the first 4 bits of the grammar. This is meant to be represent a functional pressure on the grammar.

How well a hearer’s grammar analyzes and parses utterances is measured similarly. The

combination of these measures provides an overall fitness rating.

4.3 Procedure

The basic layout of the functioning of the simulation is shown in fig. 4 as reproduced from Kirby & Hurford (1997). The diagram shows two generations of agents – on the right side triggers produced by adults are used to teach learners of the next generation. Fitness is measured based on the communicative abilities of the adults in communicative episodes between themselves (center of the diagram) and natural selection acts on the genomes of the agents illustrated on the left.

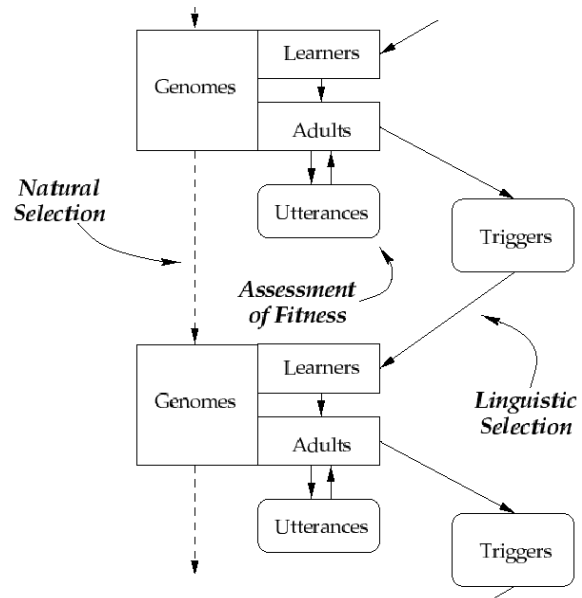


Figure 4: Procedural layout of simulation

This framework was used to test both phylogenetic functionalism and glossogenetic functionalism. Recall that a functional pressure, ease of parsability, was instantiated in the fitness function by arbitrarily introducing a bias for 1s in the first four positions of the grammar. Thus it's expected that at the end of the simulation there should be a language distribution that reflects a bias towards grammars that begin with <1 1 1 1...>. Phylogenetic functionalism would suggest that this distribution would arise solely by natural selection, whereas glossogenetic functionalism suggests that linguistic selection is *also* required.

4.4 Results

Two runs of the simulation were made, one in which a measure of parsability was included in the trigger learning algorithm (allowing for linguistic selection of parsability) and one in which the measure was not included as part of learning. In each case, the population began with 100 completely plastic agents (grammars = $\langle ? ? ? ? ? ? ? \rangle$) and the critical period was set to 200 utterances. The results were dramatically different.




Figure 6: No Linguistic Selection - Prop of alleles




Figure 5: Linguistic Selection - Prop of alleles

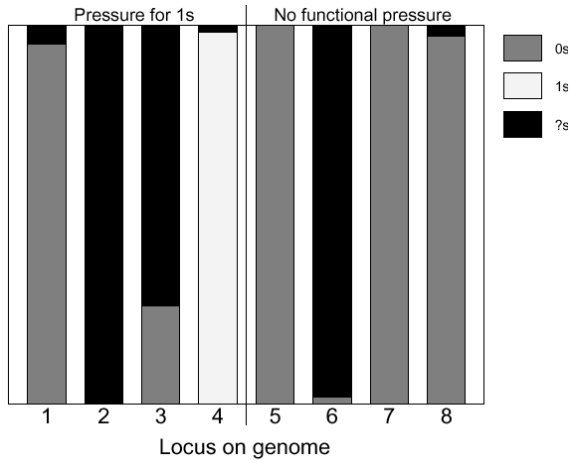


Figure 8: No Linguistic Selection - Final average

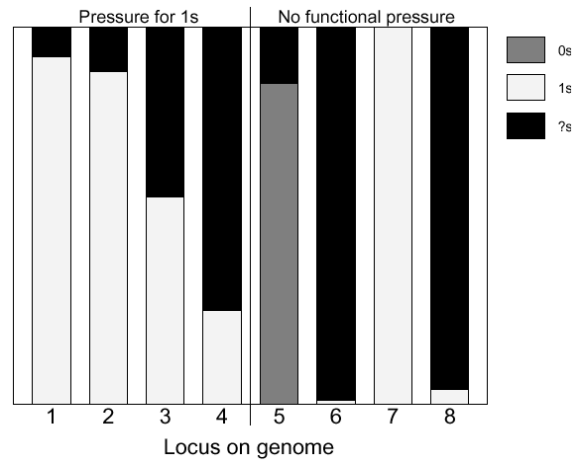


Figure 7: Linguistic Selection - Final average

Figures 5 and 6 show the proportion of 1s, 0s, and ?s in the LADs over the course of a run; 7 and 8 show the proportions of the final populations. These results are typical of multiple runs. In the first case (figs. 5 and 7) testing phylogenetic functionalism, it is clear that the simulation did not respond to the functional pressure. In fact, the first parameter has nearly completely nativised as 0, meaning that the agents in the simulation can never produce or interpret optimal utterances.⁵ On the other hand, including linguistic selection did produce the predicted result. Very quickly, in only 57 generations, the predominating the population languages (the adult grammars after learning – not the LADs) were ones that conformed to the functional pressure. Within that period only one principle had been nativised, indicating that the regularity in the adult population is the primary result of *linguistic selection* based on *cultural transmission*. However, over time the population does in fact evolve to constrain the learning possibilities of agents – i.e. what phylogenetic evolution would predict. However, it seems that it can not occur with out a glossogenetic evolution occurring first.

Evolution alone seems incapable of selecting an optimal LAD. Why might this be the case? Examining the selection pressures on linguistic agents provides a possibility; there may be a pressure to produce utterances that are easily interpretable, but a more overriding pressure is to learn the grammar of one's community. If an individual has an optimal LAD, but cannot be understood by its peers, no overall advantage is conferred on that individual.

These results imply that there are two stages in the evolution of innate linguistic constraints. A period in which glossogenetic evolution produces a subset of optimal languages that provide the regularity in the linguistic environment that is necessary for the following period of phylogenetic evolution to canalize those constraints.

⁵Similar results also occurred when fitness was based entirely on transmitting or receiving behavior.

4.5 Contribution

In a very real way, this model puts to test positions like that of Pinker & Bloom (1990), and the result supports some aspects of what they claim and does not support others. Obviously it is the case that Kirby & Hurford accept that there is an LAD which is consistent with Pinker & Bloom. But whereas Pinker & Bloom would claim that natural selection is the *only* kind of selection that can result in innate grammatical change, Kirby & Hurford attempt to not only show that natural selection is not the only source but that natural selection *cannot* act alone. Thus accepting that there is an interrelation between learning and evolution, Kirby & Hurford have provided a better understanding about the dynamics of that interrelation, showing that the assumptions that Pinker & Bloom feel are intuitively obvious, may not be correct.

The contribution this model has to the study of language evolution, should hopefully be clear. It bespeaks a complex relationship that cannot be accounted for either by nativists nor developmentalists, and it invites extensions of the model to explore the relationship in more depth.

5 Conclusion

Through these 3 models of the evolution of communication and language, I have attempted to show not that they accurately or conclusively explain the origins of human language, but more that they are a useful tool in guiding our research, testing our assumptions, and understanding the complexity of the factors involved. Contrary to detractors such as Lewontin, examining the evolution of cognition and language need not be a fruitless, baseless task. We can use modeling as a method of providing substance to our claims and upgrade our “storytelling” to “hypothesis-telling”.

Naturally, the next step is to push the models to their limits by varying their architecture, their

parameters, and their complexity. The extent to which the restrictions and characteristics found in models reflect reality and the extent to which reality is used to guide the creation of models is the extent to which we can surmise that the models support independent explanations of reality. And since we *are* basing models on what we can perceive of reality, much like evolution itself, we can expect that the models, through multitudes of generations and instantiations, will eventually display regularities, constraints, and limitations that are reflective of those found in reality. It is these reflections from which we can extrapolate general principles, that will eventually help us come to understand the nature of mind.

References:

- Botha, R. (1997a). Neo-Darwinian accounts of the evolution of language: 1. Questions about their explanatory focus. *Language and Communication* 17:3, 249-267.
- Botha, R. (1997b). Neo-Darwinian accounts of the evolution of language: 2. Questions about complex design. *Language and Communication* 17:4, 319-340.
- Botha, R. (1998a). Neo-Darwinian accounts of the evolution of language: 3. Questions about their evidential bases, logic and rhetoric. *Language and Communication* 18, 17-47.
- Botha, R. (1998b). Neo-Darwinian accounts of the evolution of language: 4. Questions about their comparative merit. *Language and Communication* 18, 227-249.
- Botha, R. (2000). Discussing the evolution of the assorted beasts called *language*. *Language and Communication* 20, 149-160.
- Deacon, T. (1997). *The Symbolic Species*. Norton: New York, NY.
- Donald, M. (1993). Précis of Origins of the modern mind: Three stages in the evolution of culture and cognition. *Behavioral and Brain Sciences* 16, 737-791.
- Elman, J., E. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett. (1996). *Rethinking Innateness*. MIT Press: Cambridge, MA.
- Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin and Use*. Praeger: New York, NY.
- Cook, V. and M. Newson. (1996). *Chomsky's Universal Grammar*. Blackwell: Cambridge, MA.
- Gazzaniga, M., R. Ivry, and G. Magnun. (1998). *Cognitive Neuroscience*. W.W. Norton & Company Inc: New York, NY.
- Gopnick, M. (1990). Feature-blind grammar and dysphasia. *Nature*, 344(6268), 715.
- Kirby, S., and J. Hurford. (1997). Learning, culture, and evolution in the origin of linguistic constraints. in P. Husbands and H. Harvey (eds.) *Fourth European Conference in Artificial Life*. 493-502. MIT Press: Cambridge, MA.
- Lewontin, R. (1998). The evolution of cognition: questions we will never answer. In D. Scarborough, and S. Sternberg (eds.) *An Invitation of Cognitive Science Vol 4: Methods, Models, and Conceptual Issues*. MIT Press: Cambridge, MA.
- Pinker, S., and P. Bloom. (1990). Natural language and natural selection. *Behavioral and Brain Sciences* 13, 707-784.
- Thelen, E., and L. Smith (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press: Cambridge, MA.