SIMULATING THE EVOLUTION OF COMBINATORIAL PHONOLOGY

W. GARRETT MITCHENER

For the workshop *Models of language evolution: Does the math add up?* at Evolang 2010.

1. Background and problem statement

The position paper by de Boer and Zuidema discusses how mathematical models and simulations are being used to develop and test theories about how the human language faculty evolved. They point out that there are significant gaps in the research community's work in this area. Primarily, there seems to be a need for modeling of biological machinery at intermediate levels of abstraction, and consequently, new simulation and modeling tools at those levels. I will discuss some of my simulations designed to fill some of these gaps. This project is in its early stages, but the preliminary results are promising.

As the position paper explains, certain parts of the overall research program in language evolution have met with great success, assisted by simulations and mathematical models. As a prime example, Nowak and his collaborators (Nowak, Plotkin, & Jansen, 2000; Plotkin & Nowak, 2000; Nowak, Krakauer, & Dress, 1999) develop an information-theoretical constraint on the lexicon and on phonology: In flat phonology, each signal requires a distinct phoneme. In combinatorial phonology, each signal is represented by several phonemes in sequence, or simultaneous phonemes transmitted over parallel channels. If there are too many topics, there may not be enough room in the space of phonemes to assign a single phoneme to each topic without incurring excessive misunderstanding. Thus, there is clear motivation for the transition from a flat phonology to an combinatorial phonology as the number of topics of conversation increases. However, details of the transition are unknown, such as what brain modifications are required to implement combinatorial signaling and what genetic mutations are required to make those modifications. This is a case where there are accurate and realistic models of the physics of speech production, and abstract models of the evolutionary forces, but I am not aware of any intermediate models that bridge the two.

It what follows, I describe a sequence of simulations of increasing complexity that bridge some of the gaps between highly abstract mathematics, physical models, and thought experiments.

2. Descriptions of simulations

Consider a problem of taking a set of Boolean inputs and producing outputs given by a fixed function of those outputs. To state the problem biologically, consider an organism that takes various actions in the presence of chemical or physical signals in its environment, thereby gaining some benefit. Let us first consider a virtual machine that represents a simplified simulation of a reaction network of genes and proteins. Complex molecules are abstracted to bit patterns. Inputs take the form of the presence or absence of specific bit patters, corresponding to signaling molecules such as the neurotransmitters and ions released by a sense organ. Outputs are given by high or low concentrations of other specific patterns. Each gene can increase or decrease the activity level of a pattern, modeling protein production, and a gene operates only when there is a sufficient amount of the protein that matches its switch pattern. The virtual machine goes through many steps in which active genes operate on pattern counts, thereby activating some genes and deactivating others, and computes a final output.

As an initial experiment, each simulated genome is considered in isolation. The genome is presented with a set of inputs, and given points for each correct output, yielding its payoff. A selection-mutation loop then steers the population toward genomes that correctly compute the desired function. To summarize the initial runs of this simulation, certain operations, such as copying inputs from one pattern to another, are easily evolved. Logical negation and compound operations require more generations and larger genomes.

Now consider the problem of transmitting information from one creature to another. To model the evolution of combinatorial phonology, suppose that the outputs of one virtual machine are transmitted across narrow channels to the inputs of another. That is, imagine that one set of inputs encodes a meaning to be transmitted, the first machine's outputs are something like motor control signals to a vocal tract, and the other set of inputs represents recognition events from a sense organ to be processed by the listening creature. We now present these creatures with the problem of copying a set of input bits across the channel. If there are few input bits, they can be transmitted directly, but if there are many, the creatures will have to develop a code and spread the message over time. The initial runs of this simulation show that this problem is distinctly harder for the artificial world to solve, however, it is able to evolve creatures that can transmit four bits across a two bit channel over time.

The next step is to first evolve creatures that must transmit few bits across a channel wide enough to accommodate them all at once, then modify the virtual world such that ever more bits must be transmitted over the same channel. This should give insights into how combinatorial phonology develops out of a fixed-size signaling system.

The overall complexity of the simulation can be increased in stages. In the

first stage, the sending and receiving creatures are identical, made from of the same genome. In the second stage, creatures from genomes selected at random are paired up for the communication task and share the payoff of success. In a later stage, each creature will be allowed to overhear conversations between other creatures to have the opportunity to learn the communal code. If at all possible, that learning process should not be built into the simulation, but rather be left to evolve. In a third stage, the channel can be based on physical models of speaking and hearing, themselves parameterized by the genome. This will lead to a bridge between Nowak et al's abstract result, and anatomical phonetics.

Eventually, I would like to use a single genome to build creatures consisting of networks of these virtual machines that each receive part of the input. Then they act together to transmit it to a receiving creature, which uses its own network to interpret the message. In this form, the virtual creatures will be complete assemblies of communication devices and accompanying behavior, evolved together.

3. Why use this approach

Other researchers have studied similar simulations, notably Cangelosi and Parisi (Parisi & Cangelosi, 2002; Cangelosi & Parisi, 1998). In their mushroom world simulations, creatures are given tasks related to identifying and processing food. Those creatures evolve the ability to complete their tasks based on signals from other creatures as well as direct information about mushrooms. They evolve compositional processing in the sense of dimension reduction of semantics: Given many bits of information about a mushroom, one speaking creature can transmit a few bits to a hearer, thereby informing the hearer of the mushroom's type. That is, the creatures evolve a system for decomposing the set of possible mushrooms into a direct product of a few features. To give another example, de Boer (Boer, 2002) describes a population evolving a discrete vowel system on historical time scales, but the vowels are represented by static formants rather than signals over time, and the main results are about how phonemes spread themselves out through learning and population dynamics.

In contrast, my project approaches the evolution of compositionality from a phonological direction. My simulation is concerned with transmitting many bits given simultaneously in their entirety over a narrow channel. Instead of semantic dimension reduction, its problem of is one of evolving the serialization and deserialization system to transmit and receive a message. The messages themselves have no meaning at this stage other than that correct transmission leads to a payoff.

The traditional artificial neural networks (ANNs) used in the mushroom world do not operate with a sense of time, so my simulation must be based on a more complicated virtual machine that does. Furthermore, my virtual machine is more similar to a biological neuron, with capabilities such as pulse generation and synchrony that are absent from traditional ANNs, so it might give greater insight into what structures to look for in the human brain that might represent composition.

If the network part of the project can be made to work, it may assist biologists in identifying how biological neurons assemble themselves as a creature grows.

My virtual machines are more time-consuming to simulate than a traditional ANN, but a single workstation can compute thousands of generations of a population of hundreds in a matter of hours or days. Thus, the increased complexity does not render the simulation computationally unfeasible.

4. Simulations and the big picture

The virtual machine I described here is designed to be more like a biological cell than is typical of such simulations. It is meant to be a tool for understanding the kinds of variation relevant to language and mental computation that mutation can discover.

Without some knowledge of the biochemistry of DNA and ontogeny, the variation available to the mutation process remains hidden, and it will be difficult to put together a realistic sequence of incremental mutations that lead to the discovery of language. Incorporating effects such as gene duplication, methylation, and nucleosome binding into an artificial life simulation involves a lot of complexity and guess work, but some attempts need to be made in this area to break the impasses discussed by de Boer and Zuidema. Likewise, the scientific community has very limited knowledge of how the brain actually represents and performs computations. It seems reasonable to simulate simplified brains of increasing realism in an attempt to aid biologists in reverse-engineering those structures and their history.

References

Boer, B. de. (2002). Evolving sound systems. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language.* Springer-Verlag.

Cangelosi, A., & Parisi, D. (1998). The emergence of a "language" in an evolving population of neural networks. *Connection Science*, *10*(2), 83–97.

Cangelosi, A., & Parisi, D. (Eds.). (2002). *Simulating the evolution of language.* Springer-Verlag.

Nowak, M. A., Krakauer, D. C., & Dress, A. (1999). An error limit for the evolution of language. *Proceedings of the Royal Society of London, Series B*, *266*, 2131–2136.

Nowak, M. A., Plotkin, J., & Jansen, V. A. A. (2000). Evolution of syntactic communication. *Nature*, *404*(6777), 495–498.

Parisi, D., & Cangelosi, A. (2002). A unified simulation scenario for language development, evolution, and historical change. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language.* Springer-Verlag.

Plotkin, J., & Nowak, M. A. (2000). Language evolution and information theory. *Journal of Theoretical Biology*, *205*, 147–159.