

Background

The evolutionary origins of the use of speech signals to refer to events or objects in the world have remained obscure. Previously, researchers interested in these questions have relied on non-human primate studies and comparative biology in order to study the origin and evolution of language. Although there have been a few exceptional cases in which chimpanzees learned to use functionally symbolic communication (Savage-Rumbaugh et al, 1980), researchers have not found actual animal languages—at least none that use basic syntax. The apparent discontinuity between human language and animal communication is problematic because, evolutionarily speaking, humans were once animals without language. It's not as though a black obelisk descended from the sky, thunder struck, and all of a sudden our ancestors had the ability to use symbols to refer to things in the environment; the human ability to use language likely developed gradually, over time, from ancestors who did not use language.

Recent advancements in computational modeling techniques have provided researchers with new tools to study the process of language evolution. Evolutionary computation techniques are particularly useful for these types of questions because they allow researchers to simulate the emergence of communication and language in evolving populations of interactive creatures (see, for example, Cangelosi, 2001). In these simulated worlds, there are hundreds of generations of autonomous agents, which are individually programmed to learn from their interactions with both objects in their environment and other agents. It is with this technique that we hope to simulate the evolution of language.

Process

The process we are going to investigate is the evolution of a communication system within a community of creatures. In our model, we will use the semiotic theory of Charles Sanders Peirce. His trichotomy of iconic, indexical, and symbolic signs has shaped the way we view the developmental stages through which communication and language evolve. Peirce's trichotomy concerns the relationship between a sign and the object it signifies. Iconic signs represent their objects through physical similarity;

indexical signs signify their objects through spatial or temporal contiguity; and symbolic signs represent their objects by arbitrary social convention.

Importantly, there is no intrinsic relationship between a symbol and its object; a symbol becomes representative of its object simply through a contract between individuals in a society. Linguistic communication systems are mostly comprised of indices and symbols, so we will focus our attention on them. In terms of communication, indexical signs point to things in the present context. For example, alarm calls are indexical because their significance is derived from co-occurring with predators in the environment. In this way, indices are strongly grounded in the environment. Symbolic signs, on the other hand, transcend the current context and environment. For example, words are mostly symbolic because they are determined purely through convention; therefore, words are meaningful on a higher level of abstraction, in which they signify concepts that are not directly tied to a particular context or environment. The only way they can be grounded is in combination with words that function indexically to specify the context. As should be apparent by now, animal communication systems generally consist of indices, while human languages operate mostly on a symbolic level.

How did the indexical communication systems of animals like primates evolve into the symbolic languages of humans? The investigation of this process is the main focus of our research, along with the original process through which an indexical lexicon (that could eventually become symbolic) might have evolved. We also hope to model the emergence of a basic syntax, which would only be possible once symbols have evolved. Overall, our goal is for our creatures to create an indexical lexicon, a symbolic communication system, and finally some simple syntax. However, we understand that these goals may be slightly far-reaching for the scope of this project, so we would feel satisfied if our creatures reached any of these stages of communication.

Because the goal of our model is to investigate a particular process—the evolution of communication—rather than the cognitive abilities of a particular species, it was not important for us to choose a specific species to focus on. As arbitrary as the choice may be, we were thinking of chimpanzees or bonobos.

Objectives

We aim to evolve language capabilities of increasing complexity: an indexical lexicon, a symbolic language, and hopefully a basic syntax (as detailed in Peircian semiotics). These are our primary goals—in contrast to the several sub-goals that define expected emergent strategies of the organisms. For example, we think the organisms will utilize language for predator avoidance (alarm calls), food quality distinctions, and expressing reproductive preferences.

Model Architecture

Our model will incorporate both evolutionary and in-life learning mechanisms of adaptation. In general, the model will use a similar evolutionary and communication system as presented in Cangelosi (2001) and associative learning model as presented in Loula et al (2003). The evolutionary model will differ from the models we have read about in two significant respects. First of all, we will enforce a genotype/phenotype distinction to make the mutation process more natural. Secondly, each generation will not be broken into rounds with an explicit selection method, but instead the creatures will reproduce and mutate in real time similar to the sheep model.

The basic cognitive model consists of the following components, ordered by activation on each tick: Saliency Evaluation, Working Memory (WM), Association Network and Conjunction Network, Behavioral Network, and finally the Behavior. The basic unit of cognition is the Predicate, a generalized (and simplified) form of C4's percepts and actions. Predicates include things like “is a predator”, “is food”, “run”, “can see”, “eat”, “breed”, “sends message 1001”, etc. Each Predicate Instance (PI) has associated with it a subject, possibly an object, and a default saliency weight. Each tick, a creature receives the true PIs of all entities near it as stimuli. That is, if it can see predator x , it will receive the PIs $Predator(x)$ and $CanSee(self, x)$. Each PI will then enter Saliency Evaluation. Based on the PI's saliency weight and the attention the creature is paying to the entities involved with the PI, the PI will enter the WM. The WM is inspired by Silas' system, except that there is one queue, rather than two, for all Predicates.

As previously specified, learning will be done via association. The Association Network is a

weighted, directed graph with predicates as vertices. The weights represent the level of association between two predicates and is updated as PIs enter the WM. The presence of PIs in the WM will cause the Predicates they are associated with to become activated (either by being added to the WM or through some other method), with the same subject/object as the original PI. The activated PIs can thus affect behavior.

The Conjunction Network is an undirected, weighted graph with Predicates as vertices that keeps track of the determined importance of the co-occurrence of certain Predicates. It updates as PIs are added to the WM. When an edge's weight reaches some threshold, a new a Predicate will form, composed of the conjunction of the two connected Predicates. The new Predicate is then available for association, behavioral activation, and entrance into the WM. In this way, we hope to achieve a kind of chunking.

Finally, the Behavioral Network links Predicates to Action Predicates (that is, predicates that the organism can perform, such as “move” and “eat”). Each link has a particular weight, positive or negative, through which the PIs present in the WM will sort of vote on which Action Predicates are to be carried out. These links will be determined both by genes and some reinforcement process not yet decided upon.

The language system incorporates naturally into this architecture. As in Cangelosi, messages consist of binary strings and have no predetermined meaning. There will be a Predicate for each possible message. That is, when one organism o sends the message 0101, all conspecifics within a particular radius will receive the PI 0101(o). Meaning can then develop through the Association Network.

The virtual world will consist of various types of food, water, and maybe foliage to hide in. There will be some predators (of yet to be determined intelligence) that serve as selective pressures.

Challenges

There are three main challenges that we foresee in our project. The first concerns the design of our synthetic creatures. This involves designing our creatures with a minimum amount of preprogramming in order to allow for emergence, while still providing them with a plausible level of innateness. Another core challenge lies in devising a method for robust, yet adaptive learning to occur. Lastly, a challenge that we foresee involves designing our environment to allow for dynamic selective

pressures.

Division of Labor

Bryan Head- Algorithm and data structure design. Chief coder. (lead guitar)

Hilary McMahan- Linguistic and semiotic hegemon; determinant for legitimacy of the model and the results through primary literature. Supporting coder. (bass & vocals)

Matthew Davis- Minister of data extraction, visualization, and statistical analysis. Supporting coder. (keyboard & rhythm guitar)

Seneca Siegel- Psychobiological legitimator of the model through primary literature and common sense. Supporting coder. (drums)

Design Review Dates

Nov. 3rd – All mechanisms will be particularized, stubbed, with associated Java documentation. Not all methods will be implemented.

Nov. 12th – A demo of the evolution systems, with some language emergence (perhaps indexical, and not symbolic communication). The learning algorithm will probably not be fully implemented, thus communication will be learned in a non-biologically relevant manner.

Nov. 24th - Entire architecture built, but parameterization will not be fully tweaked.

Dec. 3rd – Project complete and fine-tuned.

Annotated Bibliography

1) Cangelosi, A. (2001). "Evolution of Communication and Language Using Signals, Symbols, and Words." *IEEE Transactions in evolution Computation*(5): 93-101.

Cangelosi's paper deals with the question of how symbolic forms of communication evolved. He incorporates Peircian semiotic distinctions (icon, index, and symbol), Harnad's theory of mental representations for the three types of Peircian semiotics, and Deacon's hierarchy of referencing systems (based on the three levels of iconic, indexical, and symbolic relationships between objects and their representations) in order to study the development of different stages of semiotic complexity over time. The first model simulates the evolution of signal communication, ultimately finding that organisms that successfully evolve this form of language (he calls it language, I wouldn't) are better fit to survive than organisms that haven't. The second model simulates the evolution of symbol communication in a more cultural/social environment in which children imitate their parents, use parents' symbols both to decide what action to perform and to decide what name to call an object. In the evolution of symbols model, Cangelosi finds that different populations (with the same programming but with randomly generated parameter adjustment) evolved differently: some never evolved beyond single signal languages, others were capable of signal-combination, and still others showed the ability to organize their signal-combinations such that there was a parallelism between the semantic structure and the linguistic clusters (ie. they formed categories and used them in a verb-noun language).

Cangelosi's paper also lays the groundwork for other computational approaches to this subject matter. In his models, learning is brought about in two ways: agents may learn within their lifetime, through trial and error, and agents who have had more successes than failures get to contribute their 'genes' (which assigns weights to the neural net of their progeny) and thus animals are endowed with their parents' learning. Although we do not anticipate using this aspect in the model we've proposed, the methods Cangelosi uses for measuring the success or failure of his simulation are anticipated to be useful in our group's project.

2) Donaldson, M., Lachmann, M., Bergstrom, C. (2007). "The Evolution of Functionally Referential Meaning in a Structured World." *Journal of Theoretical Biology*(246): 225-233.

Donaldson et al's paper presents a game-theoretic approach to the evolution of communication. In their simulation, there are more situations than there are signals, which forces the agents to refer to multiple similar instances (or equivalence classes) of a category through a single signal. They found that ESS (evolutionarily stable strategy) must be a strict Nash equilibrium (it outperforms all other strategies when playing against itself). So, for every 'game' they run in their model, the most optimal strategy is determined for the strategy the agents are running.

Although their game-theoretic approach is not of interest to our purposes, the methods that they use to abstract the evolution of communication provide a different, but valuable, contribution to the operationalization of our model of language evolution. Specifically, their environmental design creates more situations than agents have words to describe; this can be seen as a form of selective pressure that could potentially force the formation of categorical representations in their agents.

3) Loula, A., Gudwin, R., Queiroz, J., Ribeiro, S., de Araújo, I. (2003). "Synthetic Approach of Symbolic Creatures." *S.E.E.D. Journal (Semiotics, Evolution, Energy, and Development)*(Special Issue on Computational Intelligence and Semiotics).

The researchers created a computational simulation of the emergence of symbolic communication. The group has external validity, insofar that the agents modeled were East African vervet monkeys, and the constraints on the model were derived from the actual limits of these species. They utilized Peircean semiotics as a linguistic framework for the identification of the communicative ability in the real and simulated vervet monkeys.

As far as specific implementation, the ritual world consists of monkeys and predators, and cooperation is allowed among monkeys (since they predicted intra-species communication would be advantageous). Three predators (hawk, snake, leopard) and the monkeys could be learners or teachers, and trees in the world were to provide cover from predators.

This paper is useful, since it provides us with a concrete model of language emergence in a

learner-teacher scenario. The value we derive from this is specific implementation ideas for state system, action system, and appropriate constraints on the model that will facilitate, interesting but non-trivial results.

4) Savage-Rumbaugh, E., Rumbaugh, D.M., Smith, S.T., Lawson, J. (1980). "Reference: the Linguistic Essential." *Science* 210(21): 922-924.

In this experiment, Savage-Rumbaugh, Rumbaugh, Smith, and Lawson investigated the ability of chimpanzees to organize information on a symbolic level. They began by training three chimpanzees to label six objects as either "food" or "tools". Next, the chimpanzees were asked to label objects with lexigrams for "food" and "tools", and then to do the same with photographs of objects. Lastly, and most importantly, the experimenters had the chimpanzees label object lexigrams as either the "food" lexigram or the "tools" lexigram. Two of the chimpanzees succeeded on all tasks, demonstrating an ability to conceive of the lexigrams as representational symbols. The third chimpanzee, however, failed on the object-labeling task, being unable to generalize concepts of "food" and "tool", and therefore could not move on to the subsequent trials.

This experiment reached two important conclusions that are applicable to our own model: first, the ability to organize objects as being functionally similar does not necessarily imply the ability to conceptualize objects on a symbolic level; secondly, chimpanzees do seem to be able to categorize objects symbolically, with the correct training. This is useful for our model, because it demonstrates that primates can reach the level of symbolic communication, which is a goal for the creatures in our model. However, this article also shows that we need to be very careful in determining whether or not our creatures actually reach this level.

5) Slocombe, K., Zuberbuhler, K. (2005). "Functionally Referential Communication in a Chimpanzee." *Current Biology* 15: 1779-1784.

This study examines whether or not chimpanzee 'rough grunts' functioned as referential signals. In the experiment, chimpanzees were played a sequence of rough grunts that mimic the discovery of food by another group member. A focal subject was able to use the information conveyed by these calls to guide his search for food, demonstrating that the different types of grunts were meaningful to him. This

research stands in contrast to other monkey studies which suggest that ape vocalizations are not the product of their otherwise sophisticated mentality and that ape gestural communication is more informative for theories of language evolution. This study is relevant to our proposal because it gives support to the vocal rather than gestural theories of language evolution, and could be taken to suggest that the evolution of language was not discontinuous, as some researchers have argued (for example Deacon, 1996).