The evolution of learning mechanisms supporting symbolic communication

Kenny Smith (kenny@ling.ed.ac.uk)
Language Evolution and Computation Research Unit
Department of Theoretical and Applied Linguistics
University of Edinburgh
Adam Ferguson Building, 40 George Square
Edinburgh EH8 9LL

Abstract

Oliphant (1998;1999) contends that language is the only naturally-occurring, learned symbolic communication system, because only humans can accurately observe meaning during the cultural transmission of communication. This paper outlines two objections to Oliphant's argument. Firstly, it may be that only humans possess the necessary learning mechanism to support learned symbolic communication. Secondly, such learning mechanisms may be unlikely to evolve in other species, irrespective of their capacity for observing meaning.

Introduction

Language is unique among the communication systems of the natural world - it is culturally transmitted, the relationship between basic lexical tokens and their meanings is arbitrary and those basic lexical tokens are combined to form structured forms which are used to communicate complex structured meanings. How did language come to be as it is and why is it unique?

Much recent work in the field has focused on the evolution of syntactic communication. Explanations of the human capacity for syntax have placed emphasis on two contrasting adaptive processes:

Genetic adaptation of the genetically-encoded human language acquisition device to support syntactic communication due to fitness advantages offered by syntactic communication (e.g. Nowak, Plotkin and Jansen, 2000; Pinker & Bloom, 1990).

Cultural adaptation of language in favour of compositionality, due to cultural selection resulting from language learner biases during cultural transmission of communication (e.g. Batali, in press; Kirby, 2000). Such models are not primarily concerned with the origin of the language learner's biases.

These two opposing styles of explanation offer different accounts of the uniqueness of human language. Explanations viewing the language organ as adaptive would argue that our ancestors were the only species experiencing a particular set of pressures favouring natural selection for increasingly syntactic communication. Those viewing language itself as an adapting organism have two possible explanations for the uniqueness of human language - only our ancestors underwent the mutation or set of mutations which equipped them with the necessary mental apparatus, or only the unique set of circumstances experienced by our ancestors resulted in selection for the mental apparatus required for the cultural evolution of language to begin.

Recent work by Oliphant (1998;1999), building on pioneering work by Hurford (1989), focuses on the more basic issue of the emergence of arbitrary and conventionalised word meaning. Oliphant works within the cultural adaptation framework and makes two claims. Firstly, human language is the only learned symbolic communication system. Secondly, language is unique in this respect due to the human capacity to read the communicative intentions of other language users. Once this capacity is established optimal, learned symbolic communication reliably follows through cultural evolution of communication systems.

This paper is primarily concerned with the details of Oliphant's second claim, although the insights gained from this exercise are of relevance to wider questions concerning the evolution and uniqueness of language.

Oliphant's argument

Oliphant (1998) argues that human language is the only learned symbolic communication system occurring in the natural world. Briefly, if a symbol is defined as "a sign that refers to the object that it denotes in a way that is arbitrary with respect to the process of conventinalization that established it" (Oliphant, 1998) then non-human communication systems can be classified as non-symbolic (i.e. the signs in such systems are iconic, derive from intention movements or are ritualised parts of the full behaviours they represent) or as symbolic but innate and not experientially acquired (e.g. most alarm call systems).

Oliphant (1998;1999) argues that the uniqueness of human language as a learned, symbolic communication system can be explained in terms of the problem of observing meaning during cultural transmission of communication systems. In order to learn a mapping from a signal to its meaning (rather than inherit a genetically-encoded mapping, as is the case with innate communication systems) a learner must observe the meaning-signal pairs used by other individuals. While observing a signal may be fairly straightforward, observing the meaning that signal is intended to convey may not be. For non-symbolic communication systems, the problem of observing meaning is simplified because the meaning the signal is intended to convey is contained to a certain extent in the signal, due to the non-arbitrary mapping between meaning and signal. However, no such shortcut is available when learning a symbolic communication system - the relationship between meaning and signal is arbitrary. Oliphant argues that only humans have mastered the trick of guessing which meaning a signal is intended to convey during acquisition of arbitrary meaning-signal mappings - only humans can observe meaning reliably. This capacity for observing meaning is a component of what Bloom (2000) terms theory of mind.

Oliphant (1999) supports this claim with a computational model. The structure of his argument can be characterised as:

1. We will assume that observing meaning during communication system acquisition is easy.
2. If 1 is true, very simple learning mechanisms are sufficient to support learned symbolic communication.
3. Non-human species possess these simple learning mechanisms.
4. But only humans have a learned symbolic communication system. Therefore 1 must be false - observing meaning isn't easy and only humans can do it.

As this paper centers around a computational model similar to that described in Oliphant (1999) and takes issue with conclusions raised from Oliphant's model, the model and the results obtained from it are described briefly below

Oliphant's model and results

Oliphant assumes that observing meaning is easy - learners are exposed to meaning-signal pairs during acquisition of a communication system - and investigates what type of simple learning rules result in the emergence and maintenance of optimal communication systems through purely cultural processes. The model consists of three elements:

Communication systems: Unanalysed meanings and unstructured signals are modelled using unit vectors. Communication systems therefore consist of a mapping from meaning vectors to signal vectors.

Communicative agents: Individual agents are modelled using bidirectional associative networks (see Figure 1)[1]. The mapping between meanings and signals in such networks is determined by the connection weights in the network. Prior to learning the networks have uniform weights across all their connections - the agents have no preference for any particular communication system. During training these agents are presented with meaning-signal pairs and adjust their network connection weights according to a weight-update rule. To map from a meaning to a signal, or from a signal to a meaning, the appropriate input vector is presented to the network and the output unit receiving the highest activation is activated, while the other output units remain inactive. This is a winner-take-all strategy. Given multiple output units receiving equally high activation, one winner is randomly selected.

Population model: A population consists of 100 agents. At each time step one agent is selected at random and removed. A new agent is introduced and receives three exposures to the communication systems of other members of the population. The new agent adjusts their connection weights according to their weight-update rule. This replacement process is repeated at each time step. There is no selection and no genetic transmission of connection weights or weight-update rules, although each agent in the population uses the same weight-update rule selected by the experimenter.

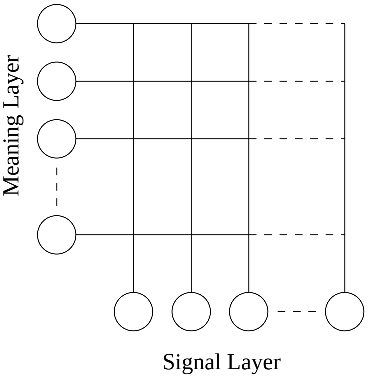

Figure 1: The associative network. Dashed lines indicate the omission of further nodes.

In this paper weight-update rules will be represented with the pattern:

$$(OnOn\ OnOff\ OffOn\ OffOff)$$

where the value in the *OnOn* slot indicates the change to be made to the connection weight between co-active units, *OnOff* is the change to be made when the meaning unit is on and the signal unit is off, *OffOn* is the reverse case and *OffOff* specifies the adjustment to make when neither meaning unit nor signal unit are active.

Oliphant defines a measure of communicative accuracy for these agents, which calculates the probability of any agent using a signal *s* to communicate a meaning *m* and having *s* interpreted as *m* by another agent. A population's communicative accuracy has a maximum of 1.0, representing the case where any meaning can be successfully communicated by any agent to any other agent in the population. Such populations are said to possess an optimal communication system.

---

[1] 10 by 10 networks are used for all models outlined in this paper.

Oliphant evaluates weight-update rules in the context of this model with respect to three criteria:

Acquisition: Can a network using the rule acquire an optimal communication system?

Maintenance: Can a network using the rule acquire an optimal system against reasonable levels of noise?

Construction: Can a network using the rule improve a non-optimal system in such a way that a population of such agents will eventually construct an optimal system?

Rules capable of acquisition will be termed *learners*. Rules capable of maintenance will be termed *maintainers* and rules capable of construction will be termed *constructors*. Oliphant evaluates three weight-update rules. The results are summarised in Table 1 (* indicates a maximum connection weight of 1).

Table 1: Oliphant's (1999) results.

Rule	Classification
(+1* 0 0 0)	learner
(+1 0 0 0)	maintainer
(+1 -1 -1 0)	constructor

Oliphant claims that the constructor rule $(+1\ -1\ -1\,0)$, a form of Hebbian learning, is well within the learning capabilities of other species. This would suggest that learned, symbolic communication systems should be common in the natural world. Oliphant argues that the rarity of such systems can therefore best be explained by the difficulty of observing meaning during cultural transmission of communication systems.

Objections to Oliphant's conclusions

This paper raises two objections to Oliphant's conclusions. Firstly, are non-humans actually capable of applying constructor-type learning procedures to the acquisition of a communication system? There is some evidence that they are not capable of doing so or simply do not do so. Secondly, if non-humans are not capable of constructor-type learning, why not? A computational model based on Oliphant's own model suggests that constructor rules may be unlikely to evolve in a population, even if members of the population are capable of observing meaning and the population is under selection pressure for communicative success.

A look at weight-update rules

A closer analysis of possible weight-update rules for associative networks suggests that the biases of constructor weight-update rules differ crucially from those of maintainer or learner rules - constructors are biased in favour of one-to-one mappings between meanings and signals. This bias appears to be present in humans, but perhaps not in chimpanzees - Oliphant's suggestion that non-human animals are capable of constructor-type learning may be incorrect.

The crucial bias What properties make a particular weight-update rule a constructor, rather than a maintainer or a basic learner? Judging this from Oliphant's sample of three weight-update rules is difficult. However, if we consider a larger array of weight-update rules the crucial characteristics become obvious. If we limit ourselves to weight-update rules which have three available actions - increase connection strength (+1 action), decrease connection strength (-1 action) or leave connection strength unaltered (0 action) there are $4^3 = 81$ possible weight-update rules (The 4 conditions *OnOn*, *OnOff*, *OffOn* and *OffOff* and 3 possible actions for each). 50 of these weight update rules are effectively nonsense - they are incapable of acquiring optimal systems and shall be termed *non-learners*. Experimenting with the remaining 31 rules in the context of populations of agents similar to those used by Oliphant reveals that 9 of the rules can be classified as constructors, 9 can be classified as maintainers and 13 can be classified as learners. The rules can be classified based on certain constraints on the relationships between the actions in the 4 conditions of the rule:

Constraint 1: $OnOn + OffOff > OnOff + OffOn$

Constraint 2: $OnOn > OnOff$

Constraint 3: $OffOn = OffOff$

Constraint 4: $OffOff > OffOn$

In order to acquire an optimal system, to be classified at least as a learner, a weight-update rule must satisfy Constraint 1. The 50 non-learner rules fail to satisfy this constraint and all the learners, maintainers and constructors satisfy it. This constraint makes intuitive sense - in order to learn a mapping, make more associations between units which tend to have matching activations than associations between units which tend to have conflicting activations.

In addition to satisfying Constraint 1, maintainers satisfy Constraints 2 and 3. Constructors satisfy Constraints 1, 2 and 4. Rules which are classified as basic learners, capable of acquiring an optimal system but not maintaining or constructing one, satisfy Constraint 1 but fail to satisfy the more rigorous constraints on maintainers and constructors in full. The constraints on constructors and maintainers are rather less obvious than Constraint 1, and need to be viewed in the context of communication system acquisition.

Consider a 2 by 2 associative network using the rule $(a\ b\ c\ d)$. Prior to learning, all the connection weights in this network are 0 - its weight matrix will be:

$$\begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

If this network is exposed once to $m1$, the vector $(1\ 0)$, paired with $s1$, the vector $(1\ 0)$, its weight matrix will be:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

For both constructors and maintainers Constraint 2 applies - $a > b$. This means that if our 2-D network agent uses a constructor or a maintainer rule it will correctly produce $s1$ to communicate $m1$, due to the winner-take-all output procedure.

Constructors and maintainers differ in their relationship between c and d. For constructors, Constraint 4 applies - $d > c$. In the context of our 2-D network, this means that if the network uses a constructor rule it will automatically prefer to use $s2$ to communicate $m2$, despite the fact it has only been trained to associate $m1$ with $s1$. This is the crucial property of constructors - they are biased in favour of one-to-one mappings between meanings and signals. As a result there will be strong cultural selection for one-to-one mappings between meanings and signals in populations of constructors. As the members of the population learn from each other a shared one-to-one mapping will gradually emerge, and a shared one-to-one mapping is an optimal system of communication.

Are maintainers biased in favour of one-to-one mappings? They are not - because Constraint 3 applies, $c = d$, our 2-D network using a maintainer rule will be equally likely to express $m2$ using $s1$ or $s2$, due to their equal weights in the network. Maintainers are neutral with respect to one-to-one mappings - they can learn them in the presence of noise, provided that that noise does not drown out the one-to-one mapping.

Do learners have any biases with respect to the nature of the mapping? Learner rules form a less homogeneous group than constructors or maintainers, but all learner rules are biased against one-to-one mappings and in favour of many-to-one mappings from meanings to signals. This explains why they cannot maintain an optimal system in the presence of noise - strong cultural selection in populations of learner agents results in any noise being amplified and spread until the population converges on a mapping where all meanings are expressed by one signal which is therefore minimally informative.

Human and chimpanzee learning biases Oliphant contends that constructor-type learning mechanisms are widespread in the natural world. But is there evidence that any species is actually biased in favour of learning one-to-one mappings, particularly in the domain of communication? Humans appear to be. It has been suggested that word acquisition in humans is guided by the Contrast Principle (Clark, 1988), a bias in favour of one-to-one mappings between meanings and words. While Bloom (2000) suggests this principle is part of the human theory of mind, it can be conceived of as a communication-specific learning bias. Do any other animals have this kind of bias?

There is some evidence from ape language learning research that chimpanzees, our closest extant relatives, do not posses this one-to-one learning bias and are in fact biased in favour of many-to-one mappings. Kanzi, a pygmy chimpanzee, “has used *clover* to refer to the specific plant, but he also uses it to refer to parsley that grows in tight clusters on the ground and to red bud blossoms that grow in tight cloverlike clusters on a tree”(Savage-Rumbaugh et al, 1986). Kanzi also “uses a food name, such as *juice*, to indicate that he wants to go to the location ... where juice is typically found” (Savage-Rumbaugh et al, 1986). This many-to-one mapping between meanings and signals could be attributed to some kind of similarity-based extension of signals. More persuasive evidence for the lack of a Contrast Principle in chimpanzees comes from Kanzi’s younger sister, Mulika, and his mother, Matata. During the process of acquiring the lexigram-based communication system, Mulika “began by using the lexigram *milk* for many different things, including requests to be picked up, requests for attention, requests to travel to different places, requests for food and requests for milk” (Savage-Rumbaugh et al, 1986). This phase lasted for 2 months. Matata “did not develop an adequate concept of one-to-one correspondence between a given symbol and a given referent” (Savage-Rumbaugh et al, 1986). This evidence suggests that chimpanzees may not be biased in favour of one-to-one mappings and may be biased in favour of many-to-one mappings - in terms of this paper, chimpanzees may not be capable of constructor-type learning, and may in fact be better characterised as learner-type learners. Contrary to Oliphant (1999), the correct learning mechanism for supporting learned symbolic communication systems may be specific to language acquisition and specific to humans.

Constructor learners fail to reliably emerge

If constructor-type learning biases are rare in the natural world, this offers another possible explanation for the uniqueness of language as the only learned symbolic communication system - only humans have evolved the necessary learning mechanisms. But why? A slightly altered version of Oliphant’s argument seems appealing - only humans can observe meaning, and once this trick has been mastered natural selection will quickly deliver up a learning bias which supports symbolic communication. But can constructor rules reliably emerge, even if we assume a pre-existing capacity for observing meaning? A computational model similar to that outlined in Oliphant (1999) suggests that such weight-update rules may be unlikely to evolve, even under apparently ideal conditions.

The new model and results

Two modifications were made to Oliphant’s (1999) model to allow the evolution of weight-update rules conducive to symbolic communication:

1. In Oliphant’s model all members of the popula-

tion used the same weight-update rule, which was selected for the population by the experimenter. In the new model, the weight-update rule used by each agent is genetically specified in a 4-locus genome corresponding to the phenotype rule $(\text{OnOn OnOff OffOn OffOff})$. There are 3 possible alleles for each locus: $+1$, 0 or -1, corresponding to the 3 possible changes in connection strength in the phenotype network. This genome can therefore encode any of the 81 weight-update rules analysed above. The initial population in all simulations consists of agents with random genotypes.

2. In Oliphant's model, removal from the population (death) was random and there was no notion of breeding as all agents were identical at birth. In the new model selection pressures on death and breeding are introduced. At each time step the population are evaluated for communicative accuracy according to the measure used by Oliphant (1999)[2]. Less successful communicators are more likely to be removed from the population and more successful communicators are more likely to breed to produce new offspring, to whom they transmit their genetically-encoded weight-update rule[3]. The population is therefore under natural selection for communicative success.

In common with Oliphant's model, the agents were assumed to be able to observe meaning during the cultural transmission process. This model therefore seems perfect for evolving learning biases conducive to communicative success - agents can observe meaning, weight-update rules are genetically encoded and there is strong selection pressure for communicative accuracy. Do populations under these conditions reliably converge on constructor rules and develop optimal, learned symbolic communication?

Figure 2 illustrates the progress of a run which converged on a (very near) optimal communication system. This is a typical example of a successful run. However, only 5 of 100 runs were successful in this respect - optimal communication is unlikely to emerge, even under this ideal set of circumstances.

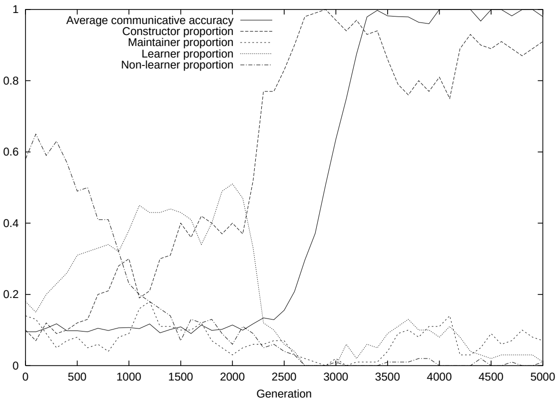

Figure 2: A successful run.

Why does natural selection have such difficulty in identifying weight-update rules which lead to optimal communication? The problem is that constructors need time to converge on an optimal system - cultural selection over repeated time steps gradually moves their communication systems into increasingly optimal overlapping areas of communication system space until they are all converged on an optimal system. In the new model, where weight-update rules are genetically transmitted, in the early stages of this construction process individuals using constructor rules have little fitness advantage over other individuals. As a consequence the genetic transmission process will be essentially random - the population will undergo genetic drift. In successful runs, such as that shown in Figure 2, genetic drift preserves constructors, by chance, in sufficient numbers for sufficient time to allow the construction process to get well under way. Constructors then show increased communicative accuracy which leads to steady selection for constructor genes, constructor numbers in the population increase and the population's average communicative accuracy increases sharply.

Interestingly, when the population's communicative accuracy nears optimal levels, selection for constructors cuts out - the population enters a second stage of genetic drift, where constructor numbers fluctuate randomly. This is due to the fact that maintainers, and to a lesser extent learners, are capable of putting the finishing touches on a communication system and maintaining that system once it is established, given a little assistance by selective removal of poorly-communicating agents. Constructors loose their fitness advantage over maintainers and learners and genetic transmission becomes semi-random once more. However, non-learners never drift back into the population - they are incapable of learning an optimal system and suffer a severe fitness penalty. This three-stage drift-selection-drift pattern is common over all successful runs and is illustrated in Figure 3.

These simulation results suggest a further objection to Oliphant's argument that the key factor in explaining the uniqueness of human language is the human capacity for observing meaning. Even if other species could observe meaning accurately the evolution of a learning mechanism which supports optimal learned symbolic

---

[2] Each individual is evaluated for their ability to successfully transmit and receive 20 randomly selected meanings with randomly selected members of the population.

[3] Tournament selection is used for both death and breeding. For death, a three-agent tournament is held, with the agent whose communicative accuracy is rated lowest being removed. For breeding, two parents are selected. Each is selected using a three-agent tournament, with the agent whose communicative accuracy is rated highest being allowed to breed. One-point crossover occurs during breeding with probability 0.95. Mutation occurs at each locus on the offspring genotype with probability 0.01. Mutation results in replacement of the allele at the mutated locus with one of the other two possible alleles, randomly selected.

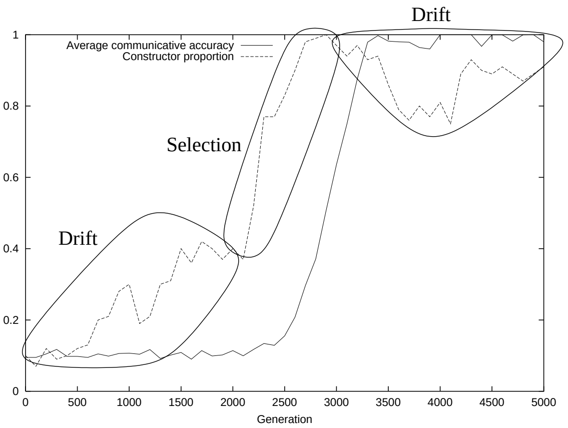

Figure 3: The three stages of a successful run.

communication would not necessarily follow. It could be that other species can observe meaning, but humans are an evolutionary fluke in hitting upon a constructor-type learning mechanism, or that the appropriate learning mechanism in humans became fixed under more reliable selection for some other useful application.

Conclusions

Oliphant (1998;1999) suggests that optimal, learned symbolic communication will trivially emerge given a capacity to observe meanings during cultural transmission of communication and a commonly-occurring learning bias, and that language is the only naturally-occurring instance of such a system because only humans have the ability to observe meanings accurately. This paper raises two objections to Oliphant's proposal:

1. The learning bias necessary to create an optimal learned communication system, a bias in favour of one-to-one mappings from meanings to signals, may be specific to humans.
2. Even if other species were capable of observing meaning, the correct learning bias would be unlikely to evolve due to the delay between its emergence and any fitness reward to those possessing it. The human learning bias may represent an evolutionary fluke or a reappropriated piece of mental apparatus.

The question of whether non-human animals possess an appropriate form of one-to-one bias could be answered by studying signal acquisition in a context equivalent to that described in this paper for a 2 by 2 network. When trained to associate $m1$ with $s1$, do non-humans exhibit a preference for $s2$ when asked to communicate $m2$, as do constructors? Or do they reliably exhibit a preference for $s1$, as some of the literature on ape language learning suggests? It may be that non-humans such as Kanzi possess a weaker one-to-one bias, acquire the one-to-one bias during development or never acquire it.

The second part of this paper, summarised in 2 above, represents an attempt to integrate the genetic adaptation and cultural adaptation accounts of the evolution of language, which must be the next goal of researchers in this field. While the arising models and their behaviour are somewhat complex, such an approach must be preferred to one which rules out a particular adaptive process as a starting assumption.

Acknowledgements

Thanks to Simon Kirby, Jim Hurford, Henry Brighton, Andrew Smith, Peter Todd and Jason Noble for comments on this and earlier versions of this work.

References

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

Bloom, P. (2000). *How Children Learn the Meaning of Words*. Cambridge, MA: MIT Press.

Clark, E.V. (1988). On the logic of contrast. *Journal of Child Language, 15,* 317–335.

Hurford, J.R. (1989). Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua*, *77*, 187–222.

Kirby, S. (2000). Syntax without Natural Selection: How compositionality emerges from vocabulary in a population of learners. In C. Knight, M. Studdert-Kennedy & J.R. Hurford (Eds.), *The Evolutionary Emergence of Language: Social function and the origins of linguistic form*. Cambridge: Cambridge University Press.

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

Oliphant, M. (1998). *Rethinking the language bottleneck: Why don't animals learn to communicate?* Unpublished manuscript, Department of Theoretical and Applied Linguistics, University of Edinburgh.

Oliphant, M. (1999). The learning barrier: Moving from innate to learned systems of communication. *Adaptive Behavior*, *7(3/4)*, 371–384.

Pinker, S. & Bloom, P,. (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, *13*, 707–784.

Savage-Rumbaugh, S., McDonald, K., Sevcik, R.A., Hopkins, W.D. & Rubert, E. (1986). Spontaneous Symbol Acquisition and Communicative Use by Pygmy Chimpanzees (Pan paniscus). *Journal of Experimental Psychology: General*, *115(3)*, 211–235.