Published in Language Learning and Development: Hauser, M., Barner, D., & O'Donnell, T. (2007). Evolutionary linguistics: A new look at an old landscape. *Language Learning and Development*, 3(2), 1-32.

Please Note: This manuscript is posted for noncommercial use. It is the pre-publication version and, as such, may contain errors that are not present in the final publication.

Evolutionary Linguistics: A New Look at an Old Landscape

Marc D. Hauser¹, David Barner^{1, 2}, and Tim O'Donnell¹

¹ Harvard University

² University of Toronto

Abstract

This article explores the evolution of language, focusing on insights derived from observations and experiments in animals, guided by current theoretical problems that were inspired by the generative theory of grammar, and carried forward in substantial ways to the present by psycholinguists working on child language acquisition. We suggest that over the past few years, there has been a shift with respect to empirical studies of animals targeting questions of language evolution. In particular, rather than focus exclusively on the ways in which animals communicate, either naturally or by means of artificially acquired symbol systems, more recent work has focused on the underlying computational mechanisms subserving the language faculty and the ability of nonhuman animals to acquire these in some form. This shift in emphasis has brought biologists studying animals in closer contact with linguists studying the formal aspects of language, and has opened the door to a new line of empirical inquiry that we label *evolvingo*. Here we review some of the exciting new findings in the *evolvingo* area, focusing in particular on aspects of semantics and syntax. With respect to semantics, we suggest that some of the apparently distinctive and uniquely linguistic conceptual distinctions may have their origins in nonlinguistic conceptual representations; as one example, we present data on nonhuman primates and their capacity to represent a singular–plural distinction in the absence of language. With respect to syntax, we focus on both statistical and rule-based problems, especially the most recent attempts to explore different layers within the Chomsky hierarchy; here, we discuss work on tamarins and starlings, highlighting differences in the patterns of results as well as differences in methodology that speak to potential issues of learnability. We conclude by highlighting some of the exciting questions that lie ahead, as well as some of the methodological challenges that face both comparative and developmental studies of language evolution.

Introduction

If one looks back at the history of research into the evolution of language, it is possible to detect at least three distinct phases of exploration. Phase 1 can be dated back to the ancients, with continued analysis until the late 19th century. Its signature was speculation, often with interesting Gedanken experiments such as the rearing of children on an island without exposure to language. Its theoretically unconstrained approach was so significant that the Linguistic Society of Paris deemed, on March 8, 1866, that it would not allow any discussion concerning “l’origine du langage, soit la creation d’une langue universelle.” That is, it was strictly verboten to speak about either the origins of language, or the possibility of a universal language. The advisory board had clearly reached its limit on speculation, and would wait until someone had something empirical to contribute.

Phase 2 hailed the arrival of a highly empirical period, accompanied by new theoretical insights. Beginning around 1950–1960, significant contributions emerged from two distinctively different traditions. On the one hand, scientists trained in comparative psychology initiated research into the possibility of teaching apes spoken English or some other human-created symbol system. The goal was to assess whether other animals could learn the tricks of our linguistic trade. Attempts to teach chimpanzees to speak resulted in utter failures. Looking to the comparative anatomical work of Negus (1929, 1949) would have provided the essential clue: chimpanzees, along with other nonhuman primates, do not have the kind of vocal tract that could, even with proper motivation, generate the sounds of human speech (Fitch, 2000). Inspired by these failures, several research teams explored the possibility that chimpanzees, gorillas, and orangutans, as well as dolphins, might acquire visual symbol systems that correspond in important ways to human signed and spoken words. These studies were far more successful, and

continue into the present. At about the same time as the training studies started, ethologists kicked off an exploration of the natural communication of animals from a broad comparative perspective, including most notably, von Frisch's (1967) work on the honeybee dance language. The aim of this work was to assess whether the spontaneously generated communicative signals of animals showed some of the same design features as human language.

Unlike the ban that emerged in the 19th century, the second phase was met with excitement, as well as a healthy dose of skepticism and criticism (Wallman, 1992). For those excited by the discovery of apes using American Sign Language (ASL) or artificially created symbols that mapped onto rich conceptual structures, there was hope that we might one day close the gap on our understanding of language evolution (Gardner&Gardner, 1969; Premack, 1971). If apes, and some other animals, could acquire a symbol system with the expressive power of human words, then our uniqueness would be challenged, and we would have a story beautifully consistent with Darwin's theory of gradual evolution by natural selection. Similarly, if some of the design features of human language were also present in animals (Hockett, 1960a, 1960b), then we could speak about precursors to language in the same way that we speak of precursors to our opposable thumbs.

Although reviewing in detail the history of research in Phase 2 is beyond the scope of this article (for synthesis of this work, see Bradbury & Vehrencamp, 1998; Cheney&Seyfarth, 1990; Hauser, 1996), we would like to offer one diagnosis. Though we have learned a great deal from this research, and continue to do so today, our general sense is that much of it has had little impact on the fields of linguistics and psycholinguistics. We realize that this seems like an awfully dismissive, critical, and pessimistic conclusion. We offer it here as only an observation based both on a reading of recent research in these areas, as well as on the basis of a comparison

with what we consider Phase 3. We therefore make a few points about Phase 2 to transition into Phase 3.

If you look at any recent book in linguistics or psycholinguistics, in any of the primary areas of research including phonology, semantics, and syntax, you will barely see a citation to research on animals from Phase 2. If the work is mentioned, it is typically dismissed as being irrelevant or methodologically flawed. In some cases, we think the critiques have been unfair or premature. The training studies have come under attack from at least two different perspectives. On the one hand, many have charged that the work is irrelevant because of training. Humans don't acquire language from hours and hours of training and because the process is so very different, the work on apes is simply not relevant to those studying humans. We personally think that this charge misses the point in at least one critical way: if the question at hand is whether nonhuman animals can evidence certain core computational or perceptual processes involved in language, we can explore this problem with methods that involve either spontaneous responses or training. What is critical is to distinguish between questions concerning the uniqueness of certain human capacities as opposed to how they are acquired. If animals show evidence of a capacity, but require training to express it, then this provides an example of a species difference in learning as opposed to a difference in the target capacity per se. In contrast, if animals are incapable of expressing the capacity even under a training regime, then the capacity is arguably unique to humans. Further, and foreshadowing issues to come, if an animal has a capacity that can only be demonstrated by means of training under artificial conditions, then this raises serious questions, often noted by Chomsky, concerning its possible evolutionary function. On the one hand, it would appear to represent a genuine puzzle or paradox for biology that evolutionary processes might build structures that have no expressed function. On the other hand, the

revolution in molecular biology and development that we are witnessing today reveals numerous cases in which there are phenotypically unexpressed traits, but the genetic potential to express them. For example, recent work on insects reveals that the basic genetic building blocks for constructing wings are present in wingless insects, but not expressed (Whiting, Bradler, & Maxwell, 2003). As Whiting et al. stated, “Our results support the hypothesis that the developmental pathway for wing formation evolved only once in insect diversification, but that wings evolved many times by silencing and re-expressing this pathway in different lineages during insect evolution.”

The second level of critique has focused on the nature of the evidence itself, and the extent to which it has any bearing on the nature of human linguistic phenomena. In the case of several studies of apes, and in particular, some of Savage-Rumbaugh’s (Savage-Rumbaugh, 1986; Savage-Rumbaugh & Lewin, 1996; Savage-Rumbaugh, Murphy, Sevcik, Brakke, Williams, & Rumbaugh, 1993; Savage-Rumbaugh, Rumbaugh, Smith, & Lawson, 1980) work on bonobos, although the use of lexigrams to communicate was certainly trained, and some of the competence to comprehend spoken English may come from social reinforcement, we find it unfair to rule out this work either because of training, or because some of the abstract operators differ from those involved in language. For example, instead of nouns, verbs, and adjectives, the bonobo Kanzi seems to access agent, object, and location, although the latter set of abstractions are different from the former and may not represent the core, some, such as agent, certainly play a central role in almost all theories of syntax, especially as a feature that helps define verb classes. What does seem fair, however, is that the vocabulary that these animals have acquired appears to be based on a cognitive process that is fundamentally different from that underlying human words. Specifically, though these animals may acquire some 300 to 500 distinctive

symbols, and can use them in the appropriate context, their lexicon never takes off as it does in very young children, nor is there any sense in which they grasp the combinatorial or compositional nature of these symbols; rarely do animals combine two or more symbols, and when they do, the system appears highly constrained (Kako, 1999). In this sense, linguists and psycholinguists consider the work to provide little traction with respect to either questions of evolution or perhaps more important, the extent to which our linguistic capacity relies on unique resources.

The work on animal communication received a similar response from linguists and psycholinguists. Though the pioneering studies of vervet monkey alarm calls (Seyfarth, Cheney, & Marler, 1980; Struhsaker, 1967), together with follow-up studies of vervets and other animals calling in the same and different contexts (Arnold & Zuberbuhler, 2006; Dittus, 1984; Gouzoules, Gouzoules, & Marler, 1984; Hauser, 1998; Kiriazis & Slobodchikoff, 2006; Marler, Dufty, & Pickert, 1986; Zuberbuhler, Cheney, & Seyfarth, 1999), provided gorgeous evidence for highly communicative signals, these systems differ from human words on many levels (see also, Tomasello, this issue). In parallel with the critique of the signing studies, none of these natural communication systems appear open ended with respect to their expressive power. If animals had access to a system that linked sound arbitrarily to meaning, then there would be no constraints on their repertoire size; instead of the 20 to 25 different calls observed in most ape and monkey species, we should see hundreds if not thousands of calls, or a smaller set that can be combined into an infinite variety of strings. Given that we don't, the system must be quite different from the one underlying the system of semantics in human language. Further, with the exception of the honey bee's dance language, all of the calls discussed above, and produced by birds and mammals, indicate objects and events in the here and now, with no evidence of signals

indicating past experiences or ones anticipated. Finally, there is no evidence that these calls fit into abstract categories approximating nouns, verbs, or adjectives, and no evidence that animals take into account the mental states of their audience when conveying such information (see also Tomasello, this issue). For these reasons and others, the work has contributed little to our understanding of how humans acquired the capacity to link particular chronological representations with particular conceptual representations to create words, either signed or spoken, with the power to refer to mind internal and external properties of the word, both abstract and varying in their temporal dimensionality.

We turn next to Phase 3, which is only a few years old, but we believe indicative of a new period of growth, associated with a more amicable and productive collaborative effort between biologists, psychologists, and linguists (Christiansen & Kirby, 2003; Hauser, Chomsky, & Fitch, 2002; Hurford, Studdert-Kennedy, & Knight, 1998; Jackendoff, 1999; Pinker & Jackendoff, 2005; Pullum & Rogers, 2006). We introduce Phase 3 in the next section, associating this period with the birth of a new discipline that we call *evolutionary linguistics*, ” or *evolvingo* for short.

The Aims Of Evolvingo

Though not explicitly stated at the time, the ideas put forward by Hauser, Chomsky, and Fitch (2002) and subsequently clarified by Fitch, Hauser, and Chomsky (2005) provide the general aims of evolvingo. A significant aspect of this approach is a distinction between language qua internal computational capacity of the mind and language qua communication (see also Tomasello, this issue). We favor the idea that these two should be kept separate, and both studied. It is of course true that we use the internal computations of the mind to communicate, generating an expressed language. But it is also true that we often don’t convey our internal

thoughts and even when we do communicate with language, the mapping between internal and expressed will be imperfect. The reason for keeping these processes separate is not only theoretically useful, but methodologically important. As discussed above, the history of research on language evolution has focused almost entirely on the communicative aspects of our capacity for language, as opposed to the internal computations that enter into language comprehension and production. As a result, when students of animal behavior have contributed to the ongoing debates about language evolution, they have done so by focusing on communication, either spontaneous or trained; two exceptions to this general pattern were Premack (Premack, 1986; Premack & Premack, 1983, 2002) and Pepperberg (Pepperberg, 1987a, 1987b, 2000), who focused more squarely on the conceptual foundations of language, their presence and absence in chimpanzees and African gray parrots respectively. By separating the internal computations from what they might or might not be used for, we open the door to a new line of inquiry, and especially, the possibility that computations deployed by humans for language may be deployed by animals in other domains such as spatial navigation and social interaction that have no communicative expression. Thus, for example, though animals as different as insects, fish, birds, and mammals use dead reckoning and landmarks to navigate, none of these taxonomic groups, with the exception of insects, and honeybees in particular, convey to others where they have been or where they are going. There is thus a dissociation between their own internal representations and their capacity to express these in communication. But the fact that these internal representations are abstract, and often provide optimal solutions to adaptive problems, makes their exploration crucial in the context of potentially isomorphic representations in humans, including those underlying language.

Keeping the distinction between internal computation and communication in mind, Hauser et al. 2002 proposed a distinction between the faculty of language in the broad (FLB) and faculty of language in the narrow sense (FLN). FLB is simply those processes of the mind that are both necessary and sufficient to support language. Thus, for example, attention is involved in language processing but is neither unique to language nor unique to humans. FLN includes those processes that are both uniquely human and unique to language. Providing a demonstration that a particular process is uniquely human doesn't clinch its place in FLN, as it is possible that this uniquely human process is implemented in other domains of human knowledge, and thus, not unique to language. For example, some kind of combinatorial process generating hierarchical structure is operative in music, mathematics, and morality, making this computational operation a candidate component of FLB.

Hauser et al. 2002 hypothesized that FLN, though potentially an empty set, may only include the computational resources subserving recursion and their interface or mapping to the conceptual-intentional (semantics) and sensorimotor (phonetic) systems; though there are many different technical (and nontechnical) definitions of the term *recursion*, they all have at their core a rule, procedure, or condition that is defined in terms of itself and a so-called "base case," which ends the otherwise infinite regress. Recursion in formal systems leads to discrete infinity. On this view, FLN would include interfaces between mechanisms implementing recursion and the systems of phonology and semantics; the Faculty of Music in the Narrow sense (FMN) might include the interface between recursive operations and the systems of pitch perception and emotion. The key programmatic issue then is that our understanding of language evolution relies critically on a broad comparative exploration, covering different species, different methods, and different domains of knowledge. It is this programmatic feature of Hauser et al. 2002 that we

stress. We turn next to a brief discussion of why the comparative angle is important, both taxonomically and methodologically.

The Importance Of A Broad Comparative Perspective

The FLB–FLN distinction forces a phylogenetic question for any given Capacity X: Is it uniquely human? There are two ways to think about the answer to this question. The simplest is this: finding any species with Capacity X rules out the claim that it is uniquely human. More interesting, for some, is evidence that a monkey or ape has Capacity X. This kind of evidence, more narrow in taxonomic scope, is potentially more interesting because it raises the distinct possibility that we acquired X by common descent from an ancestor that also had Capacity X; that is, if both human and nonhuman primates have Capacity X, the similarity is most likely a homology as opposed to an analogy. Both homologies and analogies are relevant to questions of uniqueness, and both inform different aspects of the problem. Homologies reveal evidence of common descent, whereas analogies (often) reveal evidence of convergence and specifically, the possibility that common selective pressures have led to the appearance of the same trait in two distantly related species. Thus, for example, the neural and psychological mechanisms subserving face processing in human and nonhuman primates are extremely similar, and most assume that this is evidence of homology. In contrast, the mechanism subserving the acquisition of song in Passerine songbirds and language in humans is similar, but most likely an example of analogy given that nonhuman primates show very different patterns of acquisition (Egnor & Hauser, 2004).

Though much of the comparative work on language evolution has focused on nonhuman primates, we hope that this taxonomic myopia will soon change. In particular, even if we fully test the primate order, and find no evidence that humans and nonhuman primates share a set of

computational capacities, we have not provided a satisfactory test of FLN. In fact, we haven't even capitalized on the logic of Darwin's comparative method. The goal isn't to mindlessly test every species under the sun, but rather, to think about the ways in which even distantly related species might share common ecological or social problems, thereby generating common selective pressures and ultimately, solutions given a set of constraints. Further, once we liberate ourselves from a narrow focus on communication, we can begin to look at the possibility that some of the computations subserving language in humans are actually deployed in other domains in both humans and other animals, and thus, did not evolve for language (Hauser, 2006).

The Importance of a Broad Methodological Perspective

Tied to our emphasis on a broad taxonomic scope is an emphasis on a broad methodological scope. And again, there are both theoretical and methodological reasons for this move. Theoretically, we need to test animals with a variety of methods because a single method can never generate firm conclusions about cognitive capacities. This conclusion is true in a trivial sense, but also in a much more interesting sense. The trivial sense follows from the fact that if we test an animal on a given task and it fails, we have gained virtually no understanding of why. For any given failure, there are a multitude of reasons for the failure. If they succeed, it is then important to zoom in on reasons for success, and here is where the theoretical issues become of interest. Assume, for the sake of argument, that we have trained an animal in a classic operant task, and following training, the animal shows that it can solve a particular problem. This shows that the animal has the capacity to solve the problem. But what we don't learn from this experiment is whether this is the only way that the animal can evidence such abilities. For example, is operant training or shaping necessary? Could the animal demonstrate competence at the task without shaping? These questions are important given the comparative issues at hand

(Macphail, 1987).

To clarify the comparative point, let's return to language. Many of the interesting phenomena we see in human language production and comprehension arise spontaneously, without training. Take the capacity to use and comprehend words (Bloom, 2000; Pinker, 1994; Tomasello, 2005). At some point in development, usually around the first birthday, children spontaneously begin to use words, often comprehending a far greater number beforehand. Rapidly, their lexicon grows, reaching some 30,000 to 50,000 words by high school. Tallying up the sounds produced naturally by animals, or the symbols acquired by apes and dolphins, the numbers range from about 20 to 400. There are two things to note here. First, in the case of natural vocal signals, the numbers are extremely small relative to human language. Second, even in the case of training, they are also small. What this suggests is something important about both the capacity and its acquisition: nonhuman animals do not acquire anything like the capacity to refer with words that we do, neither spontaneously nor following massive training. This is only one case, developed more fully in the next section. The important point is that we should be prepared to observe three distinctive patterns when we carry out the full comparative approach, contrasting our capacity for language with that of other animals:

1. Humans are uniquely endowed with capacities not found in other animals.
2. Humans and other animals share some of the same capacities, but humans acquire them spontaneously whereas animals require training.
3. Humans and other animals share some of the same capacities, and acquire them spontaneously.

All three patterns are interesting in terms of comparative distribution as well as the interface between computational mechanisms and acquisition devices.

The Evolution of our Semantic System

Without conceptual content and linguistic meaning, there would be little to talk about, and thus little reason to link representations of FLN to sensorimotor systems for producing and processing speech sounds. It is therefore crucial to understand how conceptual representations are linked to other systems of language, and whether over the course of linguistic evolution certain concepts became available to humans that were not previously available to other animals. Did the evolution of language also involve the evolution of language-specific concepts?

As noted earlier, much early research on animals focused on the capacity to comprehend symbolic or referential relations between signs, vocalizations, or external symbols and objects in the world. This research program produced many fascinating results, but above all suggested that the human capacity for acquiring a large, practically unbounded lexicon is probably not shared with other species. Further, the way in which other animals use the symbols or signs that they do acquire suggests that for them, symbolic meanings are not compositional: the novel combination of two familiar symbols does not result in the automatic assignment of a novel meaning to the combination, as is often the case for human language. Though researchers have actively searched for syntax in animal calls, the results of such studies have revealed nothing like the capacity expressed by natural language users. For example, Arnold and Zuberbuhler (2006) recently reported that pudge-nosed monkeys string together two different call types—one that draws attention to leopards (call it “A”) and one that signals eagles (call it “B”)—to create a communicative signal with a different meaning than either call on its own (“A + B” meaning “let’s move on”). For this to be of relevance to language (again, as opposed to communication more generally), one would need evidence that the animals’ use of these particular sounds exhibits certain properties. To be syntactic, the A + B call would need to be derived from the

calls A and B via a structural operation or rule (because the A + B pairing could potentially arise without reference to the A and B calls individually). To conclude that this syntax resembles that of natural language (i.e., exhibiting compositionality), we would need evidence that the meanings of each individual call combine to produce a new meaning (i.e., that A + B is derived from A and B by a syntactic operation, and that this operation has the potential to take other inputs and yield different, and novel, results (Fodor, 1998; Fodor & Lepore, 2002; Kamp & Partee, 1995; Partee, 1995).¹ Short of evidence that animals employ a logic of this kind, it is difficult to conclude that the calls are structured and interpreted according to a syntactic system like that found in natural language.

Why do animals fail to spontaneously communicate their mental states using symbolic systems, even when these symbols are provided over years of training? As research on animal conceptual understanding progresses, it becomes increasingly apparent that what nonhuman animals lack is not concepts to express (Cheney & Seyfarth, 1998; Hauser, 2000), but rather a mechanism for combining concepts and for outputting them to systems that can externalize them (e.g., permitting expression via sign or sound, see Tomasello, this issue). An important body of

¹ The evidence for structurally dependent compositionality in language is captured well by the example of subsecutive adjectives like big and tall (see (Barner & Snedeker, 2007; Kamp & Partee, 1995; Partee, 1995) for experimental evidence of compositional knowledge in young children). When subsecutive adjectives like big (A) combine with nouns like fish (B) to form a noun phrase like big fish (A + B), actual big fish are identified according to the typical sizes of fish (Bierwisch, 1989; Cresswell, 1976; Kennedy, 1999; Klein, 1991). Fish that are bigger than average, or bigger than most other fish, are categorized as big fish. The range of sizes specified by big thus depends on the noun that it combines with and the set of things in the world this noun is used to pick out (contrast big mouse vs. big elephant). For this reason, inferences of the following kind are not valid (Partee, 1995):

(1) Goldie is a big fish.

Goldie is a pet.

Therefore Goldie is a big pet.

As recently shown by Barner and Snedeker (2007), even 4-year-old children are able to use this type of compositional knowledge to apply adjectives like tall to novel sets of objects, and to shift what counts as tall depending on the typical sizes of kind members.

research has examined categorization and conceptual understanding in animals in domains like number, object tracking, and theory of mind and navigation, among others (Cheney & Seyfarth, 1990; Gallistel, 1990; Hauser, 2000; Shettleworth, 1998; Tomasello & Call, 1997). Studies of these capacities suggest that nonhuman animals have a rich repertoire of conceptual representations that organize their experience of the world. These studies have often paralleled investigations of human infants, sometimes finding large overlap and shared mechanisms across species. In many domains, animals appear to have as much to talk about as human infants, but unlike humans, do not develop a way of putting these ideas together and outputting novel combinations to external symbolic systems. Even animals highly trained in sign or using symbols fall far short of expressing the full extent of their knowledge in any given domain and fail to spontaneously communicate their internal states via a symbolic system.

One interesting and understudied piece of this puzzle is function words and morphology in natural language, which express set relations, quantification, tense, aspect, and more. This functional morphosyntax depends almost entirely on other words for its interpretation. Also, although languages differ widely in the extent to which different functional morphemes are explicitly encoded, they are acquired with little effort when made available, and sometimes emerge spontaneously in home-signing children, who lack any relevant input whatsoever (Goldin-Meadow, Mylander, & Franklin, in press). Currently, little to nothing is known about whether the conceptual distinctions expressed by functional morphosyntax are represented prelinguistically by human infants, let alone nonhuman animals. This, we believe, is a crucial frontier for understanding the evolution of language, for it seems that if there are to be any concepts that are unique to human language, they may be those whose meanings are inextricably bound to their relations to other words, and thus are inherently compositional.

To make matters more concrete, consider an example. In languages like English, a composed meaning can be derived by combining a content word, like *apple*, with a function word or determiner, like *a* or *some*. Thus, we can refer to a single apple as an apple or a plural set of apples as some apples. Singular–plural morphosyntax combines with the count noun apple to derive two distinct composed meanings. These functional bits and pieces can combine freely with any noun, including ones that have just been acquired, to derive an unlimited number of singular and plural noun phrases. However, singular–plural morphosyntax is uninterpretable in isolation (uttering “a” or “-s” without the content word apple is meaningless, in a way that uttering apple alone is not). If animals truly lack the capacity to combine old concepts to derive new ones, then they should have no use for concepts like the singular–plural distinction, which depend on compositionality for their interpretation. However, if some animals have this capacity, but merely fail to express it using human-trained external symbol systems, then we might expect to find evidence for related conceptual distinctions in experimental paradigms that do not depend on animals’ acquisition of explicit symbol systems. We therefore suggest that an important research path for students of language evolution is not only an exploration of animals’ understanding of symbols; conceptual domains like number, space, objects, structural relations, and sensorimotor interface systems, but also an assessment of animals’ knowledge of concepts that underlie abstract morphology (e.g., function words, etc.). In the following section we flesh out this argument and provide one empirical example.

Quantificational Systems

Perhaps the most promising area to begin this investigation is natural language quantification. In languages like English, individuals, sets, and subsets can be represented via singular–plural morphology and, as mentioned above, by lexical quantifiers like *some*, *most*, and

all, and by measure phrases like *piece of* and *bit of*. The course for studying the conceptual foundations of these expressions is hopeful, because so much is already known about the range of nonlinguistic quantificational abilities in humans and nonhuman animals. Decades of research has revealed that many species are capable of representing and comparing the numerical properties of sets. This numerical capacity has been described as stemming from two distinct representational systems. First, the analog magnitude system, sometimes called the number sense, represents the approximate cardinal value of sets, and has been documented in a wide variety of species including rats, pigeons, rhesus monkeys, human adults and human infants (Cantlon & Brannon, 2006; Cordes, Gelman, & Gallistel, 2002; Dehaene, 1997; Feigenson, Dehaene, & Spelke, 2004; Flombaum, Junge, & Hauser, 2005; Gallistel, 1990; Hauser & Carey, 2003; Hauser, Tsao, Garcia, & Spelke, 2003; Whalen, Gallistel, & Gelman, 1999). Second, the system for parallel individuation is often described as a visual object tracking system, and is characterized by a sharp limit in the number of objects that can be simultaneously represented and tracked. These representations permit the comparison of two sets on the basis of one-to-one correspondence, so long as each set is within the limits of the system. For human adults, this limit is normally three to four objects, although young infants are normally limited to tracking three objects in parallel (Feigenson & Carey, 2003, 2005; Hauser, et al., 2000; Pylyshyn, 2003; Uller, Jaeger, & Guidry, 2003). Rhesus monkeys have a set limit of four objects (Hauser et al., 2000), which permits comparisons in which both sets have four members or less (e.g., two vs. four), but results in failure for comparisons where one set has more than four members (e.g., three vs. eight).

Neither of these systems alone provides representations of the kind expressed by natural language quantifiers. Linguistic set-relational representations fail to encode the absolute

magnitude of sets in the way of analog magnitudes. In addition, they have no set limit like parallel individuation, and do not track the identity of individuals within particular sets. Instead, set-relational expressions like plural nouns make reference to sets of an unspecified, and potentially unbounded magnitude, and fail to distinguish between pluralities of different sizes (e.g., some apples can refer equally to a set of 2, 3, or 10 apples). However, the ability to perform numerical comparisons using either analog magnitudes or parallel individuation does require a prior assignment of individuals to distinct sets. For example, to determine that one set contains more individuals than another requires first segregating the objects under consideration *qua* individuals into two distinct arrays, prior to assessing the approximate magnitude of each (or comparing the sets based on one-to-one correspondence). Thus, some primitive, underlying system of set representations must exist. The exciting question for *evolvingo* is whether these set representations are related to those found in human natural language. Do nonhuman animals deploy set representations that resemble those that human infants bring to the task of language acquisition?

In a recent study (Barner, Wood, Hauser, & Carey, *in review*), we began to explore this question by testing whether rhesus monkeys have representations that support a distinction between singular and plural sets, when other numerical representations (e.g., analog magnitudes and parallel individuation) are not available. Our starting point was to figure out a way of stripping down set representations to their bones, creating experimental circumstances in which other number systems would not be activated. Previous research from prelinguistic infants and nonhuman primates suggests that this is possible given two conditions: (a) the comparison is restricted to small sets (i.e., four objects or fewer), such that analog magnitude representations are not primed, and (b) individuals in each set do not move independently, but are presented as

united sets moving as a collection, such that parallel individuation (or object tracking) is not primed, and comparison cannot be based on one-to-one correspondence. Research on human adults indicates that, when objects undergo common motion, moving as collections, subjects track collections and assign a single visual index to the collection as a whole. Under such circumstances, if subjects track whole sets but not the individuals that comprise them, then parallel individuation should not allow comparison of set magnitudes using one-to-one correspondence. Also, under various experimental conditions, neither human infants nor rhesus monkeys apply analog magnitude representations to sets within the limits of parallel individuation. Infants discriminate between large sets (e.g., 4 vs. 8, 8 vs. 16) on the basis of the ratio between them, but not between small sets with equivalent ratios (e.g., 2 vs. 4; Feigenson, 2003; Feigenson & Carey, 2005; Xu, 2003; Xu & Spelke, 2000). Similarly, in the absence of training, rhesus monkeys do not spontaneously use analog magnitude representations when one of the sets being compared contains fewer than four objects, and when objects are presented sequentially (Hauser et al., 2000).

We hypothesized that when small arrays of objects are presented as united sets, neither analog magnitude nor parallel individuation representations would be engaged to represent sets. Therefore, we predicted that rhesus monkeys—the subjects of our study—would be unable to make numerical comparisons based on these systems, but would be restricted instead to using conceptual distinctions provided by set representations alone. For example, we hypothesized that without analog magnitude representations, subjects would be unable to resolve the difference between two plural sets based on their relative magnitudes. Similarly, without parallel individuation, we predicted that animals would be unable to compare the sizes of small sets (e.g., two vs. four) based on one-to-one correspondence. However, if rhesus monkeys can represent the

distinction between singular and plural sets in the absence of parallel individuation and analog magnitudes, then they should distinguish comparisons like one versus two and one versus five but fail to distinguish between two sets of more than one (e.g., two vs. four).

We tested wild rhesus monkeys on Cayo Santiago, Puerto Rico using a forced choice foraging paradigm previously used with this species and with human infants to study number representations (Feigenson, Carey, & Hauser, 2002; Hauser et al., 2000). In the first experiment, animals watched as an experimenter placed a set of apples in one box, and then a second set in another. The experimenter then moved away from the boxes, allowing the animals to choose between them. Crucially, apples in each set were attached with Velcro to foam core boards in a row (spaced 3 cm apart), such that when they were revealed and hidden in the box they always moved together as a united set. Under these conditions, animals chose the larger set for comparisons of one versus two and one versus five, but not for comparisons of two versus four and two versus five. Thus, rhesus monkeys succeeded at singular–plural comparisons, but not at plural–plural comparisons. Also, their pattern of success was based neither on parallel individuation, which predicts success involving sets of four or less (e.g., two vs. four), nor on analog magnitudes, which predicts success based on ratio (e.g., similar behavior for one vs. two and two vs. four, or any comparisons with a greater ratio, like two vs. five). In a second experiment, we showed that representations of parallel individuation can be triggered if objects are presented sequentially. Using the method of Hauser et al. (2000), we placed apples for each set in the boxes one at a time, and monkeys succeeded at one versus two and two versus four but not two versus five or three versus five. This pattern of results is predicted by parallel individuation. Further, we found success at one versus five comparisons, consistent with the idea that more basic distinctions supported by set representations themselves are preserved when

other systems, like parallel individuation, are activated.

A final set of experiments using the forced choice foraging paradigm investigated the specific conditions under which singular–plural representations are triggered. In one condition we presented rhesus monkeys with comparisons (one vs. two, one vs. five, two vs. four, and two vs. five) where apples were first seen as united sets placed together on a board and then moved one at a time into the boxes. Thus, objects were seen first as united sets but last seen as independently moving individuals. In the second condition the order of this manipulation was reversed. Apples were placed one at a time onto a presentation board and then hidden as united sets undergoing common motion into the boxes. When objects were last seen as united sets, monkeys again succeeded at one versus two and one versus five, but failed at two versus four and two versus five. Thus, they succeeded only with singular–plural comparisons. When they were last seen as independently moving individuals, monkeys succeeded at one versus two, one versus five, and two versus four but not two versus five. Thus, we replicated the findings of the first experiment, suggesting that animals use parallel individuation only when objects are last seen as independently moving individuals.

To summarize, this series of studies provides evidence that rhesus monkeys have the capacity to distinguish between two sets without representing their absolute magnitudes and without comparing them on the basis of one-to-one correspondence. Small sets of objects failed to trigger analog magnitude representations under all conditions, and parallel individuation was activated only when objects underwent independent motion. Further, we found that when neither parallel individuation nor analog magnitudes were activated, set representations supported a distinction between singular and plural sets, regardless of the absolute magnitude of the plural sets.

These initial studies provide hope for investigating the evolutionary origins of more subtle conceptual distinctions expressed by natural language. It is a hallmark of linguistic quantification that set relations can be described without appeal to the absolute magnitude of the sets being referred to, and without indexing each individual. Expressions like some cats can refer to 2, 4, 5, or 100 cats. The acquisition of lexical quantifiers, plural morphology, and the classifiers depends on the availability of a numerical hypothesis space that is not concerned with magnitude or object tracking, but with relations between sets and individuals. Thus, the finding that exactly this type of representation can be elicited in nonhuman animals opens the door to an empirical program for studying the origins of concepts that support linguistic quantification. Needless to say, this program will not only focus on primates, but other species as well, given the strong possibility that these conceptual abilities are broadly shared, with ancient evolutionary roots.

The study of the evolutionary precursors to semantic structures offers many avenues for exploration. For one, the results from rhesus monkeys raise the question of whether similar results would be found in prelinguistic infants using appropriate measures. If such representations are indeed relevant to language acquisition, then we might expect them to be found in humans as well as in nonhuman animals. Previous studies have found signatures of parallel individuation when objects are presented as independently moving individuals, but have never systematically explored the status of these representations when objects undergo common motion. We are currently exploring this in 15-month-old infants to determine whether previous failures of infants to resolve singular–plural comparisons like one versus four (Barner, Thalwitz, Wood, Yang & Carey, in press; Feigenson & Carey, 2003, 2005) are attributable to the sequential presentation of set members. Second, we have begun investigating animals' capacity

to represent other set relations. Can animals distinguish between sets and subsets in the way that humans do when distinguishing all from some of a set? Can they represent the conceptual distinction expressed by the quantifier *most*, independent of the total number of objects presented in sets? How are set relations related to animals' representation of property and category information, such as that between mass and count nouns?

This research can be paired with continuing studies of categorization and conceptual representation to flesh out the full extent of nonlinguistic conceptual representations that might have formed the origins for natural language semantics over evolutionary history. Do animals' representations of sets resemble those deployed by natural language in humans? Does their existence require limited compositional capacities? The studies presented here suggest the possibility that animals have some compositional capacity, deployed for reasoning about sets of different kinds, but that this capacity is not available to input–output systems for generating and interpreting novel combinations of symbols. For this, humans may have developed a unique system for linking systems of conceptual competence to sensory and motor systems, with its own compositional power: natural language syntax.

The Evolution of Linguistic Computation

Much of the research discussed so far has focused on the problem of what parts of the human conceptual system are shared with other species. Do we share the same primitive representations for number? What about other primitive concepts such as color, shape, kinship, other social relations, and so forth? Linguists refer to facts such as the existence of a shared primitive concept as substantive universals. They are “substantive” because they identify the basic units, or substance, out of which a system is built. These can be basic concepts such as number or kinship, basic categories of grammar such as “noun” or “verb,” basic distinctions

among sounds such as “voiced consonants” versus “voiceless consonants,” or any other basic inventory of categories out of which we build a cognitive system.

Linguists are also interested in a different kind of universal, the so-called formal universals. These universals are not about the inventory of basic concepts that can enter into a cognitive system, but rather about the architecture of the system itself. What kinds of rules are allowed in the system? How many different levels of representation are needed? How do different levels of representation interact? What is the structure of representations at different levels? What are the processes and algorithms by which the system is learned? Formal universals are often more abstract than substantive universals in the sense that they can only be meaningfully formulated in the context of some given formalism or computational theory. They are important because they constrain (and sometimes describe) the kinds of representations and algorithms we are likely to find in some cognitive system. If we are ultimately interested in how the brain computes linguistic structures, how such a complex system could have evolved, and what its pieces are, we ultimately must investigate formal universals.

To make the distinction between substantive and formal universals more explicit, consider an example from syntax. Early generative theories of syntax consisted of two components: a phrase structure component that generated basic constituent structures and a transformational component that contained more complex structure-dependent transformation rules. The phrase structure component consisted of a set of phrase structure rules. These had the form $X \rightarrow YZ \dots$ where X , Y , Z , and so forth were grammatical categories such as noun phrase (NP) or verb phrase (VP). These rules describe the way in which phrases represent hierarchical combinations of other phrase types. For instance, many verb phrases in English consist of a verb followed by a noun phrase—the object of the verb. This can be described by the rule: $VP \rightarrow V$

NP. Linguists tried to develop sets of such rules that described the basic constituent structures allowed in various languages. In doing so they proposed hypotheses about both substantive and formal universals. For example, the set of basic allowed phrase types and categories is a substantive universal. Does every language have a noun category and a verb category? What about adverbs? On the other hand, questions about the form of the rules pertain to formal universals. For instance, it might be useful to allow rules that were sensitive to context. For instance, we might want a verb phrase to consist of a verb marked for plurality only when it appears after a plural subject. A rule for this might look like: $\text{NP(plural)} \text{ VP} \rightarrow \text{V(plural)} \text{ NP}$. This just means that the verb inside of the verb phrase must be marked as a plural when the subject that comes before the verb phrase is. Notice however, that this particular rule is no longer of the form $X \rightarrow YZ \dots$; now we have two categories on the left-hand side: $A \text{ X} \rightarrow \text{Y Z} \dots$. In fact, it turns out that allowing rules with more complex left-hand sides like this fundamentally changes the power of the rule system—now we can describe many more kinds of structure than we could previously. Questions of whether or not this descriptive power is necessary or desirable for a theory of natural language syntax are questions about formal universals.

In the case of language, there are many competing theories of grammar that differ in both substantive and formal universals. It is easy to see how animal data can help distinguish theories that target substantive universals and give evidence about their evolution. For example, suppose it can be shown that animals that share a common ancestor with humans have a primitive distinction between “noun-like” and “verb-like” concepts. Suppose we had the choice between two grammatical theories, one of which had a basic distinction between nouns and verbs and one of which treated such a distinction as derived, rather than primitive. Although the animal data cannot “prove” either theory, we would consider the first more plausible given the existence of a

similar distinction in an animal without language.

Animal data can also be useful in distinguishing between theories that differ in terms of formal universals. Consider a situation in which there are two competing theories for some linguistic phenomenon: one divides the problem into different levels of representation by proposing that there are distinct phonological, morphological, and syntactic components. Another theory posits only phonological and syntactic components, with morphology emerging as an epiphenomenon from the interaction of phonology and syntax. The first theory, morphology, will have its own substantive universals—inventories of primitive elements such as morphemes. It will also have its own formal universals—what rules and processes it is responsible for, how it maps from inputs to outputs, what its intermediate structures look like, and how is it learned. For example, there might be a primitive operation that handles morphological reduplication, taking as input a morpheme, copying it, and adding the relevant semantic features (e.g., diminutivity). This operation might have a certain signature (e.g., it can only apply to morphemes at the edge of words).

Now suppose we are able to design an experiment where cotton-top tamarins are able to discriminate between reduplicated and nonreduplicated speech stimuli but only when the stimuli are reduplicated at edges. The tamarins would seem to be using a process formally very similar to the reduplication operation in our morphological theory. This suggests that perhaps some of the machinery underlying the morphological process is evolutionarily ancient and did not need to evolve *ex nihilo* in the human lineage. These data would support the plausibility of the linguistic theory that had the proposed reduplication operator. Once again, for obvious reasons, animal data cannot prove or disprove a linguistic theory but they must be considered as a valid source of evidence and incorporated or rejected.

It is important to be precise here about the conclusions that can be drawn from results like these. Such evidence does not lead, unambiguously, to the conclusion that the animal has morphology. In fact such a statement isn't even meaningful in this context. Morphology and its basic unit, the morpheme, are defined in the context of the human linguistic system. A morpheme is characterized precisely by the way it interacts with other parts of the language system, e.g., the lexicon, morphological rules, syntactic rules, semantic interpretation, and so forth. In the example above, the basic units that the animals were able to manipulate were not morphemes in this sense, they were simply, for example, speech syllables. The claim is not that the animal has morphology but that we have detected a possible homology or analogy in a cognitive/neural process, representation, or mechanism.

An analogy with computer programming may serve to clarify the point. One could write a routine in a programming language that sorted lists. It could take as input a list to sort and a comparator function that returned which of two list elements comes first. In one program such a routine could be put to use over lists of words, in another program it could be put to use over lists of visual objects. In both cases it might have a signature. For example, it could fail for lists over a certain size, or it could return partial results after every 10 processing steps. If we find these signatures in settings with very different kinds of list elements, we can take this as evidence that the basic routine is reused, although it might be used very differently in different cases. Likewise, in the hypothetical experiment described above, we might conclude that there is evidence that the basic computations subserved by our morphological component are built on mechanisms that are conserved across evolutionary time. In the monkeys' natural ecology, such mechanisms might be put to very different uses than they are in human language. They might not even be expressed at all. Regardless, such a discovery simply suggests the possibility of a

mechanistic analogy or homology. It is not a claim that the animal has language, morphology, or even reduplication in the linguistic sense.

Once such a candidate homology or analogy is discovered, it must be evaluated using all the normal tools. Whether it is an analogy or homology can be studied by observing its presence and use in the natural ecology and in the laboratory in a variety of closely and distantly related species. Its linguistic significance can be evaluated by seeing how consistent it is with known linguistic phenomena and theories, whether it fits into the neurological data, whether there are plausible ways in which it could have been recruited for linguistic use, and so forth. In short, it can be evaluated by using all the normal scientific tools to accept or reject hypotheses.

Artificial Language Learning (ALL)

Recent years have seen a renewed and increasing interest in the use of ALL paradigms (Perruchet & Pacton, 2006) for studying questions of formal universals in language. In ALL experiments, participants are first (either passively or actively) familiarized to some set of meaningless sequences constructed to exhibit some pattern of interest. In traditional work this set of sequences was often comprised of written strings of letters or nonsense words on cards or a computer screen. In more recent work, these sets have typically been presented auditorily as sequences of recorded or synthesized nonsense syllables. A number of studies have used visual stimuli such as temporally presented sequences of shapes (Hunt & Aslin, 2001). Once the participant has been familiarized in some fashion to this set of sequences they undergo tests designed to see if they are able to extract the pattern that the set exhibited and, in some cases, generalize it to new exemplars. These experiments can be thought of as asking two questions: (a) can participants represent the pattern in question and (b) can they learn it? The important thing about these experiments is that the basic set of sequences that the participant is expected to learn

from is meaningless. Any patterns defined in such a set must be defined distributionally. Thus ALL experiments give us leverage to model formal phenomena in language and then study them, isolated from confounding factors found in normal language such as meaning, internal structure of units, or other such “substantive” issues.

Classical ALL experiments were typically performed on adult participants or children old enough to be given explicit instructions. However, much of the renewed interest in such experiments is that paradigms have been developed that can be readily adapted to infants and animals. This allows us to ask questions about formal phenomena across different stages of human development and across species. This effort is an important part of the broad comparative program in Phase 3 of *evolvingo*. In what follows we review the main results achieved so far with animals, giving background where necessary. We focus on work that we believe reflects phenomena relevant to formal universals. For readers interested in comparative work with relevance to speech perception, see Hauser & Fitch (2003) and citations therein.

The first artificial language studies relevant to formal phenomena in language concerned prosody. Ramus, Hauser, Miller, Morris, and Mehler (2000) studied the abilities of human infants and cotton-top tamarin monkeys to discriminate between languages with different types of prosody. Linguists have proposed that languages can be classified into several basic types of prosodic structure. These are known as stress-timed, mora-timed, and syllable-timed languages (Ramus, Nespor, & Mehler, 1999). The exact status of this classification system is somewhat controversial but it seems to capture several important prosody-affecting characteristics of languages, such as whether or not they allow complex consonant clusters or vowel reduction. Ramus et al. showed that both human infants and cotton-top tamarins could distinguish Japanese and Dutch (which are very different in terms of the above classification) when stimuli were

normalized to eliminate lexical and phonetic information. However, the effect disappeared when stimuli were played backward. These results have also been replicated with rats, using operant training methods as opposed to spontaneous techniques (Toro, Trobalon, & Sebastian-Galles, 2003, 2005).

Although it is not entirely clear what features of the stimuli infants, tamarins, and rats are using to distinguish Dutch from Japanese, whatever the feature may be it is based on distributional patterns in the stimuli—a formal phenomenon. Whatever computational mechanisms are used to extract the relevant features, they are not available when the stimuli are played backward. Given that this pattern holds across rats, tamarins, and human infants, the results suggest that at least some aspects of phonetic and phonological structure are attributable to evolutionarily ancient mechanisms, perhaps explainable by properties of the mammalian auditory system.

An important application of ALL in recent years has been the study of what is known as statistical learning. Statistical learning experiments study the ability of participants to track statistical properties of the sequences to which they have been exposed. Typically these have taken the form of so-called *ngram* statistics. An *ngram* is a length N sequence of symbols (words, phonemes, etc.). For instance, a bigram could be a pair of words. Ngram statistics can take the form of frequencies (i.e., how often does AB occur in a sample), joint probabilities (what percentage of all length two sequences are specifically AB), or conditional probabilities (what percentage of the time will the presence of A predict a following of B, sometimes called transition probabilities). Computational linguistics, natural language processing, and connectionist modeling have shown that ngram statistics contain much information about the structure of natural language, and can be used for a variety of tasks such as part of speech

tagging (Manning & Schutze, 1999).

One application of particular interest is word segmentation. The naturally occurring speech stream does not contain many overt clues to word boundaries. Thus, identifying word boundaries is a major task of language acquisition for infants. One early proposal on how they might do this, originally due to Harris (1955), was that they might be sensitive to the transition statistics of syllables or phonemes in natural language. The intuition was that transitions between syllables at word boundaries would have higher entropies—more uncertainty—than transitions within words. For example, suppose that you heard the word “information” begin some sentence. There are many possible next words e.g., “is hard to come by,” “about the new product,” and so forth. On the other hand, at the point where you have heard “informa” you would probably be reasonably confident that the next syllable would be “tion.” There are not that many other possible completions of this sequence of syllables in English. If infants can keep track of the transition statistics in the speech they hear—perhaps in the form of bigram conditional probabilities— they might be able to use it to learn where one word ends and another begins in their language.

Saffran, Aslin, and Newport (1996) studied this question. They exposed infants to a continuous stream of syllables in which there were four nonsense “words.” Each word consisted of three syllables and each syllable only occurred in one word; as a result, the probability of the next syllable was always 1.0 between syllables within a word. On the other hand, words were concatenated in such a way that the probability of some given next syllable at word boundaries was always .5.

The stream was constructed so as to have no other clues to word boundaries other than the statistics just described. Saffran and colleagues showed that after listening to the continuous

stream of syllables for about 21 min, infants could reliably distinguish the three-syllable words (e.g., ABC) from both nonwords—three-syllable sequences that did not appear in the exposure (e.g., CBB)—and part-words, three-syllable sequences that spanned a word boundary (e.g., BCA). Hauser, Newport, and Aslin (2001) showed that cotton-top tamarins could also perform the same task with the same materials, reliably distinguishing words from both nonwords and part-words; and paralleling the results on prosody, Toro and Trobalon (2005) showed that following operant training, rats could also segment the continuous speech stream into distinctive words. These results suggest that the ability to learn certain kinds of ngram patterns over auditory stimuli is evolutionarily ancient.

Although English words typically have statistical structure that allows them to be identified by adjacent transitional probabilities, other languages also have words with nonadjacent structure. For example, Turkish exhibits a phenomenon known as vowel harmony where vowels in a word agree on particular phonetic features (e.g., frontness, roundness). In Turkish, if a root vowel is front (e.g., /i/), then vowels in affixes must also be front (e.g., /i/, or /e/). Semitic languages such as Arabic and Hebrew have morphology based on what is known as a trilateral root system. Vowels are inserted into three consonant long root “frames” to form different words. For example, the Arabic root “ktb” means “write,” “kataba” means “he wrote,” “maktuub” means “written,” and “makataba” means “office.” More rarely in natural language, words contain dependencies between nonadjacent syllables. For example, the Philippine language Tagalog can change verbal aspect by the insertion of the infix “-um” after the first syllable onset.

If learners could track statistics between nonadjacent consonants, vowels, or syllables, they might be able to use this information to facilitate word learning. Newport and Aslin (2004)

studied this question in human adults. They asked participants to distinguish between words and part-words in three conditions. In the first, the learners were exposed to a sequence of three-syllable words where the first syllable predicted the third syllable with perfect accuracy—that is, words had the structure 1–X–3, where X varied such that the immediate transitional probabilities were not useful in word learning. In the second condition, words consisted of three-syllable long consonant frames of the form $C_1v C_2v C_3v$. Vowels were inserted into these frames such that the transitional probability between each syllable and vowel was .5, whereas the probability between consonants was 1.0. Thus if participants were able to track consonant–consonant dependencies, as appears in Arabic, they should be able to distinguish words from part-words. The third condition paralleled the second, with the exception that the vowel–vowel transitional probabilities were now 1.0, whereas the consonants varied within the frame: $cV_1cV_2cV_3$. Human adults were able to distinguish words versus part-words in the nonadjacent consonant and vowel conditions but not in the nonadjacent syllable conditions.

Given that the distinctions between syllables and segments or phonemes and nonadjacent dependencies are core features of natural language, one might predict a priori a difference in how humans as opposed to other animals parse the continuous speech streams described above. To test this possibility, one that bears directly on the FLB–FLN distinction, Newport, Hauser, and Aslin (2004) conducted a study of cotton-top tamarin monkeys using identical materials. The pattern of results differed interestingly from that with humans: tamarins were able to discriminate in the nonadjacent syllable and vowel conditions, but not in the nonadjacent consonant condition.

The observed species difference in these nonadjacent conditional probability conditions raises several interesting questions. Given that tamarins succeed with syllables and vowels, but

not consonants, it is possible that they lack the ability to extract the relevant acoustic features that enable consonantal discrimination. Although studies using operant conditioning reveal that nonhuman animals can discriminate consonants, the problem is not in the raw perceptual mechanisms; that is, though nonhuman animals can be trained to hear the difference between consonants, in a task such as the one presented to tamarins, they appear not to do so spontaneously. Whether tamarins would do so with a different set of consonants than those used in the present experiments, or whether other animals would do so, remains an open question, as does the issue of performance by human infants. If the tamarins' failure on the consonant condition is robust, then the success on syllables can be interpreted as a success with the information in each syllable provided by the pure or co-articulated vowel nucleus. The main point here is that the comparative differences observed highlight an important point: the mechanism that computes conditional probabilities in humans and tamarins may be the same, however it may get different inputs in tamarins and humans (Hauser, Chomsky, et al., 2002). Put another way, the interface conditions between computational systems in the two species are different, due to distinctive tuning of the human sensorimotor system for consonant sounds. This fact, as opposed to a fundamentally different processing mechanism, accounts for the different pattern between species.

All the studies discussed thus far fall under the general header of statistical learning and the data can be accounted for if participants are able to track simple statistics of the input such as adjacent and nonadjacent ngrams. Doing this requires two underlying abilities. First, participants must be able to build a representation of the input that at the very least can state constraints over which words come immediately before or after (perhaps with some intervening material) another word. A system able to do this can be naturally modeled with a class of generative grammars

knows as Strictly Local systems (Pullum & Rogers, 2006). These grammars are far less powerful than even the well-known finite state grammars; given how well accepted it is in linguistics and computer science that the latter are insufficiently expressive to describe natural language, we can conclude that the former are as well.

Second, participants must have the ability to count or otherwise match probabilities in the input and attach these to their internal representations. In the jargon of computer science, they must be able to “decorate” the constraints on word transitions with probabilities. It has long been known that humans possess powerful domain-general mechanisms for these kinds of tasks. In short, therefore, the kinds of representations implied by statistical learning experiments are far too simple to be the whole story on natural language, and conversely, the mechanisms required to match probabilities are known to be domain general. Several attempts have therefore been made to test more realistic or complex structures in ALL experiments, which we turn to next.

Marcus, Vijayan, Bandi Rao, and Vishton (1999) investigated the ability of infants to acquire what they termed “algebraic rules.” Algebraic in this sense means that the rule can contain a variable—a placeholder that can allow for generalization beyond the context in which the rule was acquired. Language contains many examples of such rules. For example, in English any noun phrase can be combined as a subject with any verb phrase to form a sentence provided that they agree in number (and they themselves are otherwise well-formed phrases). So we can combine “these apples” and “taste good” as easily as “Strawberries” and “taste good.” This even applies to words we have never encountered before “blackberries taste good.” The crucial point is that verb phrase and noun phrase are typed variables, and the rule that relates them in a sentence will accept any phrase matching the type. Statistical learning as discussed above cannot account for these kinds of rules because it cannot generalize beyond the exposure sequences. To learn

algebraic rules, the learner must possess more powerful generalization mechanisms.

Marcus et al. (1999) studied the ability of 7-month-old infants to distinguish sequences of the form different-same-same (DSS) from sequences of the form same-same-different (SSD).² Infants were familiarized to three-syllable-long sequences where the second and third syllables were identical. After familiarization, infants were tested on a set of sequences half of which matched this DSS pattern and half of which matched an SSD pattern. Of note, an entirely new set of syllables was used for test items. The infants were able to discriminate the two kinds of tests; this implies that they learned the abstract pattern—the second two syllables were identical and the first was different. They could not solve the task by tracking any statistics of the actual set of syllables used in the exposure.

Given that the Marcus et al. (1999) study generated significant controversy in terms of its claims (Marcus, 1999a, 1999b), especially the idea that infants are innately endowed with the capacity to extract algebraic rules, Hauser, Weiss, and Marcus (2002) ran a virtually identical experiment with cotton-top tamarins, asking whether they too could spontaneously discriminate between DSS and SSD sequences in a habituation- discriminations task. Results showed that following habituation to DSS or SSD material, participants were more likely to respond to test material presenting a different pattern from habituation. Thus it seems that learning algebraic rules (at least of the level of complexity required by the task at hand) is not a uniquely human capacity; it is part of FLB.

Although the ability to generalize algebraic rules is undoubtedly crucial to learning natural language, language does not admit just any kind of rule. Much of modern linguistics can

² Actually, this condition was Experiment 3 of the Marcus, Vijayan, Rao, and Vishton (1999) study. However, this experiment captures the essentials of the study most cleanly and was the one used to test tamarin monkeys, so we focus on it here.

be thought of as trying to characterize the precise kinds of rules that constitute universal grammar. The Marcus experiments only represent a particularly simple kind of rule. To distinguish DSS from SSD, the learner must be able to represent at least two fundamental properties. First, the learner must represent abstract identity and its negation. That is, he or she must be able to tell if two syllables are the same or different. And second, he or she must be able to represent ordering information about syllables. He or she must be able to tell when syllables are the same or different if they came before or after other syllables (and perhaps precisely where before or after). The participant had to create a rule with three ordered variables with the restriction that the second two variables were identical whereas the first was different.

Rules in natural language, however, are often much more complicated. The arguments that fill variables can vary in length, they can be discontinuous, and they can satisfy more than one variable at a time. There are precise and complex conditions governing exactly where and when such phenomena occur. Any deep understanding of homologies or analogies with linguistic computation must take into account not just the existence of algebraic rules, but the precise kinds of rules that we find in language.

Fitch and Hauser (2004) studied questions of formal complexity of this kind. In the 1950s, Chomsky (1957) studied several systems of rules with different levels of expressivity known as phrase-structure grammars. He showed mathematically that these systems could be arranged in a hierarchy of strictly increasing generative power. Rule systems higher in the hierarchy could define all the structures of those below and other new ones.

Toward the bottom of the Chomsky hierarchy is the distinction between finite-state and higher order phrase structure grammars—the weakest of which are known as context-free grammars. Finite-state grammars can define rule systems where any sentence structure can be

computed with a finite amount of memory. Higher order phrase structure grammars require unbounded memory—the memory requirements can increase with the size of the sentence structure to be computed.

In the branch of mathematics that studies such systems—formal language theory—differences among rules are usually phrased in terms of differences in the sets of sequences (languages) that these systems can be used to define. A classic example of a finite-state language is the set $(AB)^n$. This notation just means the set of all the sequences that consist of repeats of AB (e.g., AB, ABAB, ABABAB, etc.). Notice that any computer that checks whether or not a sequence is in this set or not only needs to remember the last symbol. If it is an A then the next one must be a B. If it is a B then either the sequence can end or the next symbol must be an A.

The simplest set of sequences that is not finite-state is the set: $A^n B^n$. This is the set of sequences where the number of As is equal to the number of Bs (e.g., AB, AABB, AAABBB, etc.). Notice that for a computer to check whether or not a sequence is in this language, it needs to remember the number of As it has seen at the halfway point. Because there are an infinite different number of lengths of the sequence of As, the computer must, by implication, have unbounded memory.

Fitch and Hauser (2004) tested the ability of adult humans and cotton-top tamarins to spontaneously discriminate between languages of the form $A^n B^n$ and $(AB)^n$. The categories or types A and B consisted of a number of syllables in two different voices (a high female voice and low male voice). Participants were exposed to one or the other language and then tested on their ability to discriminate novel test strings that were either consistent or inconsistent with the familiarization material. Human adults were able to discriminate the two types of strings regardless of their exposure experience. In contrast, tamarins were only able to learn the $(AB)^n$

language. When exposed to the A^nB^n language they did not learn the rule. As Fitch and Hauser concluded, the tamarins' failure can be explained by various differences between the grammars, with A^nB^n pushing beyond the memorial capacities of the tamarins, or exceeding their pattern learning capacities in complexity of the relations within the sequences or between parts of the sequence. On the other hand, they also noted that the observed species difference do not end the comparative research, as further studies of other species, notably the great apes as well as songbirds, are needed, together with other methods, including especially operant training.

A recent article by Gentner, Fenn, Margoliash, and Nusbaum (2006; see Gentner, this issue, for additional details) followed up on the Fitch and Hauser (2004) study with two interesting extensions. On the one hand, they tested A^nB^n with European starlings, a songbird noted for its exceptional capacity for vocal learning, including a song structure that is remarkably complex in terms of its acoustic morphology and the apparent statistical patterning among notes within the song. On the other hand, they tested starlings with an operant procedure, beginning at first with $n = 2$, and using acoustic material from their species-typical song; that is, instead of the speech syllables that Fitch and Hauser implemented in their study of tamarins, and instead of the spontaneous methods, Gentner et al. used starling notes and operant conditioning. Results provided striking evidence that, following thousands of training trials, starlings can generalize to new material, including novel n (up to 6) and novel acoustic material (speech syllables).

Although these results, both the negative ones on tamarins and the positive ones on starlings, are interesting, they are difficult to interpret. There are a number of strategies available for the participant to distinguish these two languages that have nothing to do with the differences between finite-state and other kinds of grammars. For instance, as pointed out by Liberman (2004), even a first-order bigram approximation to the two languages would allow participants to

distinguish them. $A^n B^n$ has sequences that do not occur at all in the other language (namely AA and BB); the tamarins' failure is particularly striking given that this strategy would have been open to them. Second, the way that the original experiments were set up allowed humans, tamarins, and starlings to have acquired a simpler language $A^n B^m$ instead of $A^n B^n$; that is, the language of any number of As followed by any number of Bs. This language is finite-state. Gentner's study controlled for this in subsequent generalization tests, and there can be little doubt that human adults can acquire the harder language just by counting; this methodological point does, however, expose some of the difficulties inherent in these kinds of studies.

By far the biggest difficulty with these results is that they do not accurately capture the structure needed for natural language. In both the Fitch and Hauser (2004) and Gentner et al. (2006) studies, the categories A and B were not structured within themselves. That is, there was no particular B that corresponded to a particular A. Thus, participants could recognize $A^n B^n$ simply by counting. However, natural language does not use counting structures like this but exhibits so-called "dependency" structure. Individual As must correspond to individual Bs. Perruchet and Rey (2005) show that human adults are not able to extract such regularities when tested within a similar design.

In summary, although the area is new, there is already a considerable amount of comparative data relevant to the study of formal phenomena in language. There are several important directions for future work. First, we believe that the relation between artificial language studies and natural language must be clarified. This can happen both by clarifying the computational operations involved in both natural language and artificial language studies (Pullum & Rogers, 2006) and by finding empirical evidence that the same operations are used in both systems. Second, the relation between learning and representation in these studies needs to

be clarified. ALL experiments often are aimed at studying the representational capacities of a subject (e.g., can subjects represent certain kinds of grammatical dependencies?). However, they must always, by necessity, involve a learning component. For instance, in the A^nB^n experiments described above, was the tamarin's failure due to an inability to represent such structures or due to their inability to learn those structures given the sort of data they were exposed to? We feel that the subtle interplay between learning and representation has been underappreciated in work so far—leading to misinterpretation of results. It is crucial that we clarify the way that experiments inform both perspectives by exploring both the qualitative complexity bounds on ALL processes but also the quantitative bounds imposed by learning limitations. Finally, comparative work needs to be broadened to include both new species and new methods. The study of abstract computation holds the promise of finding homologies and analogies between superficially different computational systems in diverse species. These discoveries await only the imagination to develop the right methods and apply them to the right species.

Concluding Remarks

We have described three phases of research into the evolutionary origins of language. As we characterize it, we are currently in Phase 3, a period of exciting growth and development that sees terrific unification across classically isolated and often antagonistic disciplinary boundaries. In particular, formal linguistics and computer science are increasingly informing psychological research into the evolutionary and ontogenetic building blocks of the adult competence for language. The advantage of this multidisciplinary approach is that the actual language and methods of these disciplines are shared, and as a result, so too are common goals. We also note that there are major challenges ahead. Theoretically, many of the computations or operators invoked by syntacticians are extremely abstract and difficult to operationalize for purposes of

testing children and animals. Moreover, given the diversity of theoretical positions within the subdisciplines of syntax, semantics, phonology, and morphology, it is not entirely clear to those outside of the field of linguistics that particular operations or phenomena within language should be tackled from a comparative approach, and what kind of evidence would constitute an empirical demonstration in the absence of actual, expressed, linguistic constructions. To move forward, those studying nonlinguistic or prelinguistic creatures will have to work in close association with linguists who study the formal aspects of language in mature speaker/hearers to pinpoint not only the essential phenomena of language, but also the most appropriate and elegant experimental approaches to understanding their neural and psychological instantiation.

In addition to the close collaboration between scholars in different disciplines, it will be important for those contrasting different species to be clear about the nature of their methods and the conclusions that can be drawn from them. In the past, there have been unfortunate battles between researchers working on wild versus captive populations, between those using spontaneous methods and those using training, and those using species that are more closely related to humans as opposed to distantly related. These are not the right battles. A broad comparative approach at the level of taxonomy and methods is fundamental to understanding the evolution of language, and for that matter, any domain of knowledge. A careful taxonomic sampling will ultimately help shape our understanding of not only phylogenetic distribution, but also selective pressures. A careful methodological sampling will help shape our understanding of the interfaces between computation and learning. In some cases, for example, different species may share the same computational capacity, but show different patterns of acquisition. Ultimately, our best guess for the evolution of the language faculty is one that sees a loosely connected tapestry of abilities, with pieces of the tapestry shared across the animal kingdom, but

with only our own species weaving the pieces together—the patchwork quilt theory of language evolution. Thus, some species such as songbirds and whales will have capacities for vocal learning that match our own, and this will enable them to build a large repertoire of material; this capacity is not seen in our closest living relatives—the monkeys and apes. Songbirds and whales may also have some of the combinatorial and hierarchical capacities that humans have, and these are deployed in the course of arranging and rearranging their own phonological units. What these species lack, however, is a system that links the combinatorial and phonological machinery up with the system of semantics. And here is where the primates may be exceptionally gifted, and perhaps closer to humans. That is, monkeys and apes may have a richer system of semantic representations (i.e., combinatorial and open ended) but may have a limited capacity to express these representations due to impoverished connections with systems for communicative expression, which themselves lack rich combinatorial power. Studies reviewed here are consistent with the notion that many of the conceptual systems needed to support natural language may have been in place in other species prior to the emergence of language, and were woven together in humans by a species-specific syntactic innovation, enabling us to uniquely comprehend and express the infinite variety of meaningful sentences that is the gift of our nature.

References

- Arnold, K., & Zuberbuhler, K. (2006). Semantic combinations in primate calls. *Nature*, 441, 303.
- Barner, D., & Snedeker, J. (2007, November). *The integration of compositional semantics and statistics in adjective acquisition*. Paper presented at the 30th annual Boston University Conference on Language Development, Boston.
- Barner, D., Thalwitz, D., Wood, J., & Carey, S. (in press). Children's ability to distinguish "one" from "more than one" and the acquisition of singular-plural morpho-syntax. *Developmental Science*.
- Barner, D., Wood, J., Hauser, M. D., & Carey, S. (in review). Evidence for a non-linguistic distinction between singular and plural sets in wild rhesus monkeys. *Cognitive Psychology*.
- Bierwisch, M. (1989). *The semantics of gradation*. In M. Bierwisch & E. Lang (Eds.) Dimensional adjectives (pp. 71–261). Berlin: Springer Verlag.
- Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.
- Bradbury, J. W., & Vehrencamp, S. L. (1998). *Principles of animal communication*. Oxford, England: Blackwell.
- Cantlon, J., & Brannon, E. M. (2006). Shared system for ordering small and large numbers in monkeys and humans. *Psychological Science*, 17, 401–406.
- Cheney, D. L., & Seyfarth, R. M. (1990). *How monkeys see the world: Inside the mind of another species*. Chicago: Chicago University Press.
- Cheney, D. L., & Seyfarth, R. M. (1998). Why monkeys don't have language. In G. Petersen (Ed.), *The Tanner lectures on human values* (pp. 173–210). Salt Lake City: University of Utah Press.

- Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.
- Christiansen, M. H., & Kirby, S. (2003). *Language evolution*. New York: Oxford University Press.
- Cordes, S., Gelman, R., & Gallistel, C. R. (2002). Variability signatures distinguish verbal from nonverbal counting for both large and small numbers. *Psychological Bulletin and Review*, 8, 698–707.
- Cresswell, M. J. (1976). In B. H. Partee (Ed.) *Montague grammar* (pp. 261–292). New York: Academic Press.
- Dehaene, S. (1997). *The number sense*. Oxford, England: Oxford University Press.
- Dittus, W. P. G. (1984). Toque macaque food calls: Semantic communication concerning food distribution in the environment. *Animal Behaviour*, 32, 470–477.
- Egnor, S. E. R., & Hauser, M. D. (2004). A paradox in the evolution of primate vocal learning. *Trends in Neurosciences*, 27, 649–654.
- Feigenson, L. C., S. (2003). Tracking individuals via object-files: Evidence from infants' manual search. *Developmental Science*, 6, 568–584.
- Feigenson, L., & Carey, S. (2003). Tracking individuals via object-files: Evidence from infants' manual search. *Developmental Science*, 6, pp. 568–584.
- Feigenson, L., & Carey, S. (2005). On the limits of infants' quantification of small object arrays. *Cognition*, 97, 295–313.
- Feigenson, L., Carey, S., & Hauser, M. D. (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological Science*, 13, 150–156.
- Feigenson, L., Dehaene, S., & Spelke, E. S. (2004). Core systems of number. *Trends in Cognitive Science*, 8, 307–314.

- Fitch, W. T. (2000). The evolution of speech: A comparative review. *Trends in Cognitive Science*, 4, 258–266.
- Fitch, W. T., & Hauser, M. D. (2004). Computational constraints on syntactic processing in a nonhuman primate. *Science*, 303, 377–380.
- Fitch, W. T., Hauser, M. D., & Chomsky, N. (2005). The evolution of the language faculty: Clarifications and implications. *Cognition*, 97, 179–210.
- Flombaum, J., Junge, J. A., & Hauser, M. D. (2005). Rhesus monkeys (*Macaca mulatta*) spontaneously compute addition operations over large numbers. *Cognition*, 97, 315–325.
- Fodor, J. (1998). *Concepts: Where cognitive science went wrong*. Oxford, England: Oxford University Press.
- Fodor, J. A., & Lepore, E. (2002). *The compositionality papers*. Oxford, England: Clarendon.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gardner, R. A., & Gardner, B. T. (1969). Teaching sign language to a chimpanzee. *Science*, 165, 664–672.
- Gentner, T. Q., Fenn, K. M., Margoliash, D., & Nusbaum, H. C. (2006). Recursive syntactic Pattern learning by songbirds. *Nature*, 440.
- Goldin-Meadow, S., Mylander, C., & Franklin, A. (in press). How children make language out of gesture: morphological structure in gesture systems developed by American and Chinese deaf children. *Cognitive Psychology*.
- Gouzoules, S., Gouzoules, H., & Marler, P. (1984). Rhesus monkey (*Macaca mulatta*) screams: Representational signalling in the recruitment of agonistic aid. *Animal Behaviour*, 32, 182–193.
- Harris, Z. (1955). From phoneme to morpheme. *Language*, 31, pp. 190–122.

- Hauser, M. D. (1996). *The evolution of communication*. Cambridge, MA: MIT Press.
- Hauser, M. D. (1998). Functional referents and acoustic similarity: Field playback experiments with rhesus monkeys. *Animal Behaviour*, 55, 1647–1658.
- Hauser, M. D. (2000). *Wild minds: What animals really think*. New York: Henry Holt.
- Hauser, M. D. (2006). *Moral minds: How nature designed a universal sense of right and wrong*. New York: Harper Collins/Ecco.
- Hauser, M. D., & Carey, S. (2003). Spontaneous representations of small numbers of objects by rhesus macaques: Examinations of content and format. *Cognitive Psychology*, 47, 367–401.
- Hauser, M. D., Carey, S., & Hauser, L. B. (2000). Spontaneous number representation in semi-free-ranging rhesus monkeys. *Proceedings of the Royal Society, London*, 267, 829–833.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1554–1555.
- Hauser, M. D., & Fitch, W. T. (2003). *What are the uniquely human components of the language faculty?* In M. Christiansen & S. Kirby (Eds.) *Language evolution: State of the art* (pp. 158–181). Oxford, England: Oxford University Press.
- Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmenting a continuous acoustic speech stream: Serial learning in cotton-top tamarin monkeys. *Cognition*, 78, B53–B64.
- Hauser, M. D., Tsao, F., Garcia, P., & Spelke, E. (2003). Evolutionary foundations of number: Spontaneous representation of numerical magnitudes by cotton-top tamarins. *Proceedings of the Royal Society, London, B*, 270, 1441–1446.

- Hauser, M. D., Weiss, D., & Marcus, G. (2002). Rule learning in cotton-top tamarins. *Cognition*, 86, B15–B22.
- Hockett, C. F. (1960a). Logical considerations in the study of animal communication. In W. E. Lanyon & W. N. Tavolga (Eds.), *Animal sounds and communication* (pp. 392–430). Washington, DC: American Institute of Biological Sciences.
- Hockett, C. F. (1960b). The origin of speech. *Scientific American*, 203, 88–96.
- Hunt, R. H., & Aslin, R. N. (2001). Statistical learning in a serial reaction time task: Simultaneous extraction of multiple statistics. *Journal of Experimental Psychology: General*, 130, 658–680.
- Hurford, J., Studdert-Kennedy, M., & Knight, C. (1998). *Evolution of language*. Cambridge, England: Cambridge University Press.
- Jackendoff, R. (1999). Possible stages in the evolution of the language capacity. *Trends in Cognitive Science*, 3, 272–279.
- Kako, E. (1999). Elements of syntax in the systems of three language-trained animals. *Animal Learning and Behavior*, 27, 1–14.
- Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57, 129–191.
- Kennedy, C. (1999). Projecting the adjective: The syntax and semantics of gradability and comparison. New York: Garland Press.
- Kiriazis, J., & Slobodchikoff, C. N. (2006). Perceptual specificity in the alarm calls of Gunnison's prairie dogs. *Behavioural Processes*, 73, 29–35.
- Klein, E. (1991). Comparatives. In A. von Stechow & D. Wunderlich (Eds.) *Semantics: An international handbook of contemporary research* (pp. 673–691). Berlin: Walter de Gruyter.

- Liberman, M. (2004). *Hi lo hi lo, it's off to formal language theory we go*. Language Log, retrieved from <http://itre.cis.upenn.edu/~myl/languagelog/archives/000355.html>
- Macphail, E. (1987). The comparative psychology of intelligence. *Behavioral and Brain Sciences*, 10, 645–695.
- Manning, C., & Schutze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA: The MIT Press.
- Marcus, G. (1999a). Connectionism: with or without rules? *Trends in Cognitive Science*, 3, 168–170.
- Marcus, G. (1999b). Do infants learn grammar with algebra or statistics? Response to Seidenberg & Elman, Eimas, and Negishi. *Science*, 284, 436–437.
- Marcus, G., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, 283, 77–80.
- Marler, P., Dufty, A., & Pickert, R. (1986). Vocal communication in the domestic chicken. I. Does a sender communicate information about the quality of a food referent to a receiver? *Animal Behaviour*, 34, 188–193.
- Negus, V. E. (1929). *The mechanism of the larynx*. St. Louis, MO: Mosby.
- Negus, V. E. (1949). *The comparative anatomy and physiology of the larynx*. New York: Hafner.
- Newport, E. L., & Aslin, R. N. (2004). Learning at a distance. I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, 48, 127–162.
- Newport, E. L., Hauser, M. D., Spaepen, G., & Aslin, R. N. (2004). Learning at a distance: II. Statistical learning of non-adjacent dependencies in a non-human primate. *Cognitive Psychology*, 49, pp. 85–117.

- Partee, B. (1995). Lexical semantics and compositionality. In L. Gleitman & M. Liberman (Eds.), *Invitation to cognitive science, Part 1: Language* (2nd ed.) (311–360). Cambridge, MA: MIT Press.
- Pepperberg, I. M. (1987a). Acquisition of the same/different concept by an African Grey parrot (*Psittacus erithacus*): Learning with respect to categories of color, shape, and material. *Journal of Experimental Analysis of Behavior*, 50, 553–564.
- Pepperberg, I. M. (1987b). Evidence for conceptual quantitative abilities in the African parrot: Labeling of cardinal sets. *Ethology*, 75, 37–61.
- Pepperberg, I. M. (2000). *The Alex studies*. Cambridge, MA: Harvard University Press.
- Perruchet, P., & Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Science*, 10, 233–238.
- Perruchet, P., & Rey, A. (2005). Does the mastery of center-embedded linguistic structures distinguish humans from nonhuman primates? *Psychonomic Bulletin & Review*, 12.
- Pinker, S. (1994). *The language instinct*. New York: William Morrow.
- Pinker, S., & Jackendoff, R. (2005). The faculty of language: What's special about it? *Cognition*, 97, 179–210.
- Premack, D. (1971). Language in chimpanzees? *Science*, 172, 808–822.
- Premack, D. (1986). *Gavagai! or the future history of the animal language controversy*. Cambridge, MA: MIT Press.
- Premack, D., & Premack, A. (1983). *The mind of an ape*. New York: Norton.
- Premack, D., & Premack, A. (2002). *Original intelligence*. New York: McGraw Hill.
- Pullum, G. K., & Rogers, J. (2006). *Animal pattern-learning experiments: Some mathematical background*. Cambridge, MA: Harvard University, Radcliffe Institute.

- Pylyshyn, Z. W. (2003). *Seeing and visualizing: It's not what you think*. Cambridge, MA: The MIT Press.
- Ramus, F., Hauser, M. D., Miller, C. T., Morris, D., & Mehler, J. (2000). Language discrimination by human newborns and cotton-top tamarins. *Science*, 288, 349–351.
- Ramus, F., Nespor, M., & Mehler, J. (1999). Correlates of linguistic rhythm in the speech signal. *Cognition*, 73, pp. 265–292.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Savage-Rumbaugh, E. S. (1986). *Ape language: From conditioned response to symbol*. New York: Columbia University Press.
- Savage-Rumbaugh, E. S., & Lewin, R. (1996). *Kanzi*. New York: Basic Books.
- Savage-Rumbaugh, E. S., Murphy, J., Sevcik, R. A., Brakke, K. E., Williams, S. L., & Rumbaugh, D. M. (1993). Language comprehension in ape and child. *Monographs of the Society for Research in Child Development*, 58, 1–221.
- Savage-Rumbaugh, E. S., Rumbaugh, D. M., Smith, S. T., & Lawson, J. (1980). Reference: The linguistic essential. *Science*, 210, 922–925.
- Seyfarth, R. M., Cheney, D. L., & Marler, P. (1980). Monkey responses to three different alarm calls: Evidence of predator classification and semantic communication. *Science*, 210, 801–803.
- Shettleworth, S. (1998). *Cognition, evolution and behavior*. New York: Oxford University Press.
- Struhsaker, T. T. (1967). Auditory communication among vervet monkeys (*Cercopithecus aethiops*). In S. A. Altmann (Ed.), *Social communication among primates* (pp. 281–324). Chicago: Chicago University Press.

- Tomasello, M. (2005). *Constructing a language*. Cambridge, MA: Harvard University Press.
- Tomasello, M., & Call, J. (1997). *Primate cognition*. Oxford, England: Oxford University Press.
- Toro, J. M., & Trobalon, J. B. (2005). Statistical computations over a speech stream in a rodent. *Perception and Psychophysics*, 67, 867–875.
- Toro, J. M., Trobalon, J. B., & Sebastian-Galles, N. (2003). The use of prosodic cues in language discrimination tasks by rats. *Animal Cognition*, 6, 131–136.
- Toro, J. M., Trobalon, J. B., & Sebastian-Galles, N. (2005). Effects of backward speech and speaker variability in language discrimination by rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 31, 95–100.
- Uller, C. Jaeger, R., & Guidry, G. (2003). Salamanders go for more: Rudiments of number in an amphibian. *Developmental Science*, 6, pp. 568–584.
- von Frisch, K. (1967). *The dance language and orientation of bees*. Cambridge, MA: Belknap Press of Harvard University Press.
- Wallman, J. (1992). *Aping language*. New York: Cambridge University Press.
- Whalen, J., Gallistel, C. R., & Gelman, R. (1999). Nonverbal counting in humans: The psychophysics of number representation. *Psychological Science*, 10, 130–137.
- Whiting, M. F., Bradler, S., & Maxwell, T. (2003). Loss and recovery of wings in stick insects. *Nature*, 421, 264–267.
- Xu, F. (2003). Numerosity discrimination in infants: Evidence for two systems of representation. *Cognition*, 89, B15–B25.
- Xu, F., & Spelke, E. S. (2000). Large number discrimination in 6-month-old infants. *Cognition*, 74, B1–B11.

Zuberbuhler, K., Cheney, D. L., & Seyfarth, R. M. (1999). Conceptual semantics in a nonhuman primate. *Journal of Comparative Psychology*, 113, 33–42.