

16 Linguistic Adaptation without Linguistic Constraints: The Role of Sequential Learning in Language Evolution

MORTEN H. CHRISTIANSEN AND
MICHELLE R. ELLEFSON

16.1 Introduction

The acquisition and processing of language is governed by a number of universal constraints, many of which undoubtedly derive from innate properties of the human brain. These constraints lead to certain universal tendencies in how languages are structured and used. More generally, the constraints help explain why the languages of the world take up only a small part of the considerably larger space defined by the logically possible linguistic subpatterns. Although there is broad consensus about the existence of innate constraints on the way language is acquired and processed, there is much disagreement over whether these constraints are linguistic or cognitive in nature. Determining the nature of these constraints is important not only for theories of language acquisition and processing, but also for theories of language evolution. Indeed, these issues are theoretically intertwined because the constraints on language define the *endpoints* for evolutionary explanations: theories about *how* the constraints evolved in the hominid lineage are thus strongly determined by *what* the nature of these constraints is taken to be.

The Chomskyan approach to language suggests that the constraints on the acquisition and processing of language are linguistic, rather than cognitive, in nature. The constraints are represented in the form of a Universal Grammar (UG)—a large biological endowment of linguistic knowledge (e.g. Chomsky 1986). It is assumed that this knowledge-base is highly abstract, comprising a complex set of linguistic rules and principles that

could not be acquired from exposure to language during development. Opinions differ about how UG emerged as the endpoint of language evolution. Some researchers have suggested that it evolved through a gradual process of natural selection (e.g. Newmeyer 1991; Pinker 1994; Pinker and Bloom 1990), whereas others have argued for a sudden emergence through non-adaptationist evolutionary processes (e.g. Bickerton 1995; Piattelli-Palmarini 1989). An important point of agreement is the emphasis in their explanations of language evolution on the need for very substantial biological changes to accommodate linguistic structure.

More recently an alternative perspective is gaining ground, advocating a refocus in thinking about language evolution. Rather than concentrating on biological changes to accommodate language, this approach stresses the adaptation of linguistic structures to the biological substrate of the human brain (e.g. Batali 1998; Christiansen 1994; Christiansen and Devlin 1997; Deacon 1997; Kirby 1998, 2000, 2001). Languages are viewed as dynamic systems of communication, subject to selection pressures arising from limitations on human learning and processing. Some approaches within this framework have built in a certain amount of linguistic machinery, such as context-free grammars (Kirby 2000). In this chapter we argue that many of the constraints on linguistic adaptation derive from non-linguistic limitations on the learning and processing of hierarchically organized sequential structure. The underlying mechanisms existed prior to the appearance of language, but presumably also underwent changes after the emergence of language. However, the selection pressures are likely to have come not only from language but also from other kinds of complex hierarchical processing, such as the need for increasingly complex manual combinations following tool sophistication. Consequently, many language universals may reflect non-linguistic, cognitive constraints on learning and processing of sequential structure rather than an innate UG.

16.1.1 *Exploring Linguistic Adaptation through Artificial Language Learning*

The study of the origin and evolution of language must *necessarily* be an interdisciplinary endeavour. Only by amassing evidence from many different disciplines can theorizing about the evolution of language be sufficiently constrained to remove it from the realm of pure speculation and allow it to become an area of legitimate scientific inquiry. Fuelled by theoretical con-

straints derived from recent advances in the brain and cognitive sciences, the last decade of the twentieth century has seen a resurgence of scientific interest in the origin and evolution of language. However, direct experimentation is needed in order to go beyond existing data. Computational modelling has become the paradigm of choice for such experimentation (e.g. Batali 1998; Briscoe 2000; Christiansen and Devlin 1997; Kirby 1998, 2000, 2001). Computational models provide an important tool with which to investigate how various types of constraints may affect the evolution of language. One of the advantages of this approach is that specific constraints and/or interactions between constraints can be studied under controlled circumstances.

Here we point to Artificial Language Learning (ALL) as an additional, complementary paradigm for exploring and testing hypotheses about language evolution. Artificial language learning involves training human subjects on artificial languages with particular structural constraints, and then testing their knowledge of the language. Importantly, the ability to acquire linguistic structure can be studied independently of semantic influences. Because ALL permits researchers to investigate the language learning abilities of infants and children in a highly controlled environment, the paradigm is becoming increasingly popular as a method for studying language acquisition (for a review see Gomez and Gerken 2000). We suggest that ALL can be applied to the investigation of issues pertaining to the origin and evolution of language in much the same way as computational modelling is currently being used. One advantage of ALL over computational modelling is that it may be possible to show that specific constraints hypothesized to be important for language evolution actually affect human learning and processing. Below we demonstrate the utility of ALL as a tool for studying the evolution of language by reporting on three ALL experiments that test predictions derived from our evolutionary perspective on language.

In this chapter, we first outline our perspective on the adaptation of linguistic structure. Specifically, we suggest that 'language as an organism' (Christiansen 1994) provides a useful metaphor for understanding language evolution. The idea of linguistic adaptation has been explored previously using computational models of language evolution (e.g. Batali 1998; Kirby 1998, 2000, 2001). Here we report on three ALL studies that corroborate our approach. The first study, conducted by Christiansen, Kelly, Shillcock, and Greenfield (currently unpublished), points to an association between sequential learning and the processing of language. The second study,

by Christiansen (2000) with accompanying computational simulations by Christiansen and Devlin (1997), demonstrates how certain word-order constraints can be explained in terms of non-linguistic limitations on sequential learning. The third study, by Ellefson and Christiansen (2000), indicates that processes of linguistic adaptation may explain the emergence of (subjacency) constraints on complex question formation. Finally, we discuss the wider implications of linguistic adaptation for language evolution.

16.2 Language as an Organism

Languages exist only because humans can learn, produce, and process them. Without humans there would be no language (in the narrow sense of *human* language). It therefore makes sense to construe languages as organisms that have had to adapt themselves through natural selection to fit a particular ecological niche: the human brain (Christiansen 1994). In order for languages to 'survive', they must adapt to the properties of the human learning and processing mechanisms. This is not to say that having a language does not confer selective advantage onto humans. It seems clear that humans with superior language abilities are likely to have a selective advantage over other humans (and other organisms) with lesser communicative powers. This is an uncontroversial point, forming the basic premise of many of the adaptationist UG theories of language evolution mentioned above. However, what is often not appreciated is that the selection forces working on language to fit humans are significantly stronger than the selection pressure on humans to be able to use language. In the case of the former, a language can *only* survive if it is learnable and processable by humans. On the other hand, adaptation towards language use is merely *one out of many* selective pressures working on humans (such as, for example, being able to avoid predators and find food). Whereas humans can survive without language, the opposite is not the case. Thus, language is more likely to have adapted itself to its human hosts than the other way round. Languages that are hard for humans to learn simply die out or, more likely, do not come into existence at all.

The biological perspective on language as an adaptive system has a prominent historical pedigree. Indeed, nineteenth-century linguistics was dominated by an organistic view of language (for a review see e.g. McMahon 1994). For example, Franz Bopp, one of the founders of comparative linguistics, regarded language as an organism that could be dissected and classified

(Davies 1987). More generally, languages were viewed as having life-cycles that included birth, progressive growth, procreation, and eventually decay and death. However, the notion of evolution underlying this organistic view of language was largely pre-Darwinian. This is perhaps reflected most clearly in the writings of another influential linguist, August Schleicher. Although he explicitly emphasized the relationship between linguistics and Darwinian theory (Schleicher 1863; quoted in Percival 1987), Darwin's principles of mutation, variation, and natural selection did not enter into the theorizing about language evolution (Nerlich 1989). Instead, the evolution of language was seen in pre-Darwinian terms as the progressive growth towards attainment of perfection, followed by decay.

More recently, the biological perspective on language evolution was resurrected, within a modern Darwinian framework, by Stevick (1963), and later by Nerlich (1989). Christiansen (1994) proposed that language be viewed as a kind of beneficial parasite—a *nonobligate symbiant*—that confers some selective advantage onto its human hosts without whom it cannot survive. Building on this work, Deacon (1997) further developed the metaphor by construing language as a virus. The asymmetry in the relationship between language and its human host is underscored by the fact that the rate of linguistic change is far greater than the rate of biological change. Whereas Danish and Hindi needed less than 5,000 years to evolve from a common hypothesized proto-Indo-European ancestor into very different languages (McMahon 1994), it took our remote ancestors approximately 100,000–200,000 years to evolve from the archaic form of *Homo sapiens* into the anatomically modern form, sometimes termed *Homo sapiens sapiens* (see e.g. Corballis 1992). Consequently, it seems more plausible that the languages of the world have been closely tailored through linguistic adaptation to fit human learning, rather than the other way around. The fact that children are so successful at language learning is therefore best explained as a product of natural selection of linguistic structures, and not as the adaptation of biological structures, such as UG.

From the viewpoint of the UG approach to language, the universal constraints on the acquisition and processing of language are essentially arbitrary (e.g. Pinker and Bloom 1990). That is, given the Chomskyan perspective on language, these constraints appear arbitrary because it is possible to imagine a multitude of alternative, and equally adaptive, constraints on linguistic form. For instance, Piattelli-Palmarini (1989) contends that there are no (linguistic) reasons not to form yes–no questions by reversing

the word order of a sentence instead of the normal inversion of subject and auxiliary. On our account, however, these universal constraints are in most cases *not* arbitrary. Rather, they are determined predominately by the properties of the human learning and processing mechanisms that underlie our language capacity.¹ This can explain why we do not reverse the word order to form yes-no questions; it would put too heavy a load on memory to store a whole sentence in order to be able to reverse it.

Our perspective on language evolution also has important implications for current theories of language acquisition and processing. It suggests that many of the cognitive constraints that have shaped the evolution of language are still at play in our current language ability. If this is correct, it should be possible to uncover the source of some of the universal constraints in human performance on sequential learning tasks. In the next three sections, we show how language and sequential learning are intertwined, and how universal constraints on basic word order and complex question formation can be explained in terms of non-linguistic constraints on the learning of complex sequential structure.

16.3 Association between the Processing of Linguistic and Sequential Structure

The theory of language evolution presented here suggests that language has evolved to fit pre-existing sequential learning and processing mechanisms. This points to a strong association between the processing of sequential structure and language. A straightforward prediction from this approach is that one would expect impaired sequential learning and processing to lead to a breakdown of language. Indeed, Grossman (1980) found that Broca's aphasics, besides agrammatism, also had an additional deficit in sequentially reconstructing hierarchical tree structure models from memory. He took this as suggesting that Broca's area not only subserves syntactic speech production, but also functions as a locus for supramodal processing of hierarchically structured behaviour. Another study has suggested a similar association between language and sequential processing. Kimura (1988)

¹ Many functional and cognitive linguists also suggest that the putative innate UG constraints arise from general cognitive constraints (e.g. Givón 1998; Hawkins 1994; Lakoff 1987; Langacker 1987). Our approach distinguishes itself from these linguistic perspectives in that it emphasizes the role of sequential learning in the explanation of linguistic constraints.




FIG. 16.1 The finite-state grammar used to generate stimuli in Christiansen *et al.* (in preparation). Items are generated by following the arrows between nodes, and writing out their symbols. For example, the item ▲ ■ ★ ■ is produced by going from node 1 to node 2, then looping back once to node 2, followed by visits to nodes 4 and 6.

reported that sign aphasics often also suffer from apraxia; that is, they have additional problems with the production of novel hand and arm movements not specific to sign language.

More recently, our team has provided a more direct test of the suggested link between breakdown of language and breakdown of sequential learning. We conducted an ALL study using seven agrammatic patients and seven normal controls matched for age, socio-economic status, and spatial reasoning abilities. The subjects were trained on a match-mismatch pairing task in which they had to decide whether two consecutively presented symbol strings were the same or different. The materials consisted of symbol strings (e.g. ★ ● ■ ★ ● ■) generated by the simple finite-state grammar illustrated in Fig. 16.1. Subjects were instructed that they were participating in a memory experiment and that their knowledge of the string patterns would be tested later. After training, the subjects were then presented with novel strings, half of which were derived from the grammar and half of which were not. Subjects were told that the training strings were generated by a complex set of rules,² and asked to classify the new strings according to whether they followed these rules or not.

² The fact that we use rules and syntactic trees to characterize the language to be acquired should not be taken as suggesting that we believe that the end-product of the acquisition process is a set of rules. We merely use rules as convenient descriptive devices, approximating the particular grammatical regularities that we are considering.

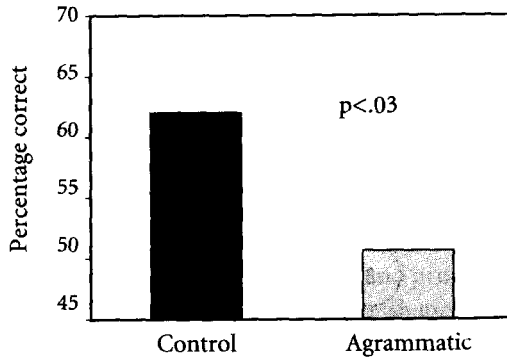


FIG. 16.2 The overall classification performance for the aphasic and normal control subjects in Christiansen *et al.* (in preparation).

The results showed that although both groups did very well on the match-mismatch pairing task, the normal controls were significantly better at classifying the new test strings in comparison with the agrammatic aphasics (see Fig. 16.2). Indeed, the aphasic patients were no better than chance at classifying the test items. Thus, this study indicates that agrammatic aphasic patients have problems with sequential learning in addition to their more obvious language deficits. This point is underscored by a recent study showing that training aphasic patients on non-linguistic hierarchical processing led to improvements in their comprehension of complex linguistic constructions (Dominey *et al.* 2001), indicating a causal link between sequential learning and language. This is, of course, what is predicted by our approach, given the suggested close connection in processing mechanisms between the learning and processing of non-linguistic sequential structure and language.

This close connection in terms of underlying brain mechanisms is further underscored by recent neuroimaging studies of ALL. Steinhauer *et al.* (2001) had subjects play a kind of board game in which two players were required to communicate via an artificial language. After substantial training, event-related potential (ERP) brainwave patterns were then recorded as the subjects were tested on grammatical and ungrammatical sentences from the language. The results showed the same frontal negativity pattern (P600) for syntactic violations in the artificial language as has been found for similar violations in natural language (e.g. Osterhout and Holcomb 1992). Another

study by Patel *et al.* (1998) further corroborates this pattern of results but with non-linguistic sequential stimuli: musical sequences with target chords either within the key of a major musical phrase or out of key. When they directly compared the ERP patterns elicited for syntactic incongruities in language with the ERP patterns elicited for incongruent out-of-key target chords, they found that the two types of sequential incongruities resulted in the same, statistically indistinguishable P600 components. In a more recent study, Maess *et al.* (2001) used magnetoencephalography (MEG) to localize the neural substrates that might be involved in the processing of musical sequences. They found that Broca's area in the left hemisphere (and the corresponding frontal area in the right hemisphere) produced significant activation when subjects listened to musical sequences that included an off-key chord. The aphasic studies and the neuroimaging studies of ALL reviewed here converge on the suggestion that the same underlying brain mechanisms are used for the learning and processing of both linguistic and non-linguistic sequential structure, and that similar constraints are imposed on both language and sequential learning. Next, we show how constraints on sequential learning may explain basic word-order universals.

16.4 Cognitive Constraints on Word Order

There is a statistical tendency across human languages to conform to a form in which the head of a phrase is consistently placed in the same position—either first or last—with respect to the remaining clause material. English is considered to be a head-first language, meaning that the head is most frequently placed first in a phrase, as when the verb is placed before the object NP³ in a transitive VP such as *eat curry*. In contrast, speakers of Hindi would say the equivalent of *curry eat*, because Hindi is a head-last language. Likewise, head-first languages tend to have *prepositions* before the NP in PPs (such as *with a fork*), whereas head-last languages tend to have *postpositions* following the NP in PPs (such as *a fork with*). Within Chomsky's (e.g. 1986) approach to language such head direction consistency has been explained in terms of an innate module known as X-bar theory which specifies constraints on the phrase structure of languages. It has further been suggested that this module emerged as a product of natural selection (Pinker 1994). As

³ NP = Noun Phrase; VP = Verb Phrase; PP = Prepositional Phrase.

such, it comes as part of the UG that every child supposedly is born with. All that remains for a child to 'learn' about this aspect of her native language is the direction (i.e. head-first or head-last) of the so-called head-parameter.

Our theory suggests an alternative explanation for word-order consistency based on non-linguistic constraints on the learning of hierarchically organized sequential structure. Christiansen and Devlin (1997) provided an analysis of word-order regularities in a recursive rule set with consistent and inconsistent ordering of the heads. A recursive rule set is a pair of rules for which the expansion of one rule (e.g. NP \rightarrow N (PP)) involves the second, and vice versa (e.g. PP \rightarrow prep NP). This analysis showed that head-order inconsistency in a recursive rule set (e.g. the rule set NP \rightarrow N (PP); PP \rightarrow NP post) creates centre-embedded constructions, whereas a consistent ordering of heads creates right-branching constructions for head-first orderings and left-branching constructions for head-last orderings (see Fig. 16.3). Centre-embeddings are difficult to process because constituents cannot be completed immediately, forcing the language processor to keep lexical material in memory until it can be discharged. For the same reason, centre-embedded structures are likely to be difficult to learn because of the distance between the material relevant for the discovery and/or reinforcement of a particular grammatical regularity. This means that recursively inconsistent rule sets are likely to be harder to learn than recursively consistent rule sets.⁴

Christiansen and Devlin (1997) also carried out connectionist simulations in which Simple Recurrent Networks (SRNs; Elman 1990, see Fig. 16.4) were trained on corpora generated by thirty-two different artificial grammars with differing degrees of head-order consistency. These networks do not have built-in linguistic biases of the sort envisioned in a UG; rather, they are biased towards the learning of complex sequential structure (e.g. Cleeremans 1993). Nevertheless, the SRNs were sensitive to the amount of head-order inconsistency found in the grammars, such that there was a strong correlation between the degree of head-order consistency of a given grammar and the degree to which the network learned to master the grammatical regularities underlying that grammar: the higher the inconsistency, the worse was the final network performance. The sequential biases of the net-

⁴ Note that our approach differs from Hawkins's (1994) performance-oriented approach to word order because he focuses exclusively on adult processing of language whereas our emphasis is on language acquisition. Although it may be impossible to tease apart the learning-based constraints from those emerging from processing, we hypothesize that basic word order may be most strongly affected by learnability constraints whereas changes in constituency relations (e.g. heavy NP-shifts) may stem from processing limitations.

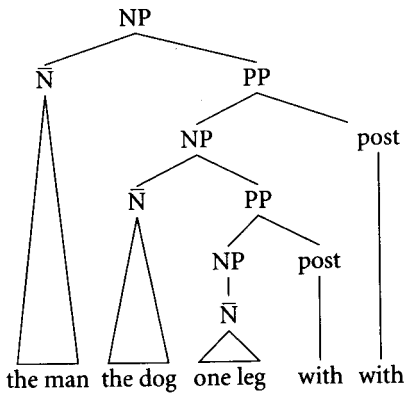
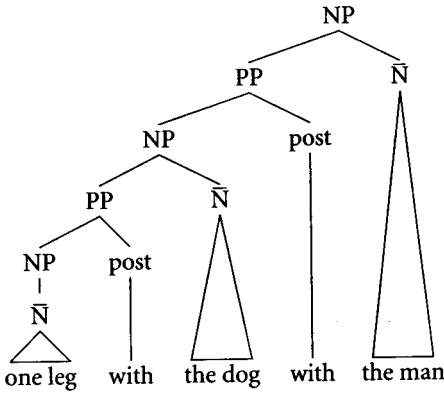
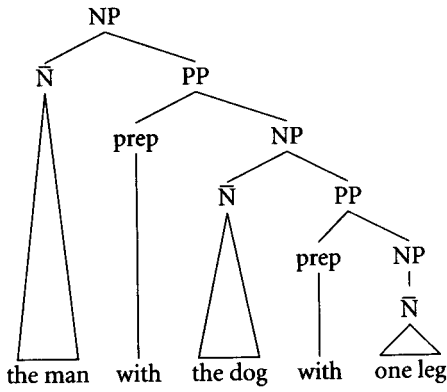


FIG. 16.3 Syntactic trees for a consistent head-first, right-branching NP (top), a consistent head-last, left-branching NP (middle), and an inconsistent NP with centre-embedding (bottom).

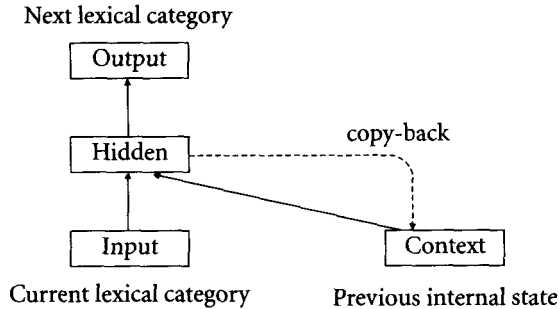


FIG. 16.4 An illustration of the SRN used in the simulations by Christiansen and Devlin (1997) and Ellefson and Christiansen (2000). The rectangles indicate sets of units, solid arrows denote trainable weights, and the dashed arrow the copy-back weights (always 1). An SRN is essentially a standard feed-forward neural network equipped with an extra layer of so-called context units. At a particular time step t , an input pattern is propagated through the hidden unit layer to the output layer. At the next time step, $t+1$, the activation of the hidden unit layer at time t is copied back to the context layer and paired with the current input. This means that the current state of the hidden units can influence the processing of subsequent inputs, providing a limited ability to deal with sequentially presented input incorporating hierarchical structure.

works made the corpora generated by consistent grammars considerably easier to acquire than the corpora generated from inconsistent grammars. This sequential learnability difference is, *ceteris paribus*,⁵ likely to result in different frequency distributions across languages through the adaptation of linguistic structure, a suggestion supported by computational simulations in Kirby (1998), showing how consistent grammars, because of their relative ease of parsing, are selected over inconsistent grammars in linguistic adaptation.

Typological analyses by Christiansen and Devlin using the FANAL database (Dryer 1992) with information regarding 625 of the world's languages further corroborated this account. Languages incorporating fragments that the networks found hard to learn tended to be less well attested than lan-

⁵ Of course, other factors are likely to play a role in whether or not a given language may be learnable. For example, the presence of concord morphology may help overcome some sequential learning difficulties as demonstrated by an ALL experiment by Morgan *et al.* (1987). None the less, sequential learning difficulties are hypothesized to be strong predictors of frequency in the absence of such ameliorating factors.

guages the network learned more easily. This suggests that constraints on basic word order may derive from non-linguistic constraints on the learning and processing of complex sequential structure, thus obviating the need for an innate X-bar module to explain such word-order universals. Grammatical constructions incorporating a high degree of head-order inconsistency are simply too hard to learn and will therefore tend to disappear. Similar simulations by Van Everbroek (1999) further substantiate this link between sequential learnability of a linguistic fragment and its frequency of occurrence. A variation on the SRN was trained on example sentences from forty-two artificial languages, varying in three dimensions: word order (e.g. subject-verb-object), nominal marking (accusative v. ergative), and verbal marking. The networks easily processed language types that occur with medium to high frequency amongst the languages of the world, while low frequency language types resulted in poor performance. Together, the simulations by Christiansen and Devlin and Van Everbroek support a connection between the distribution of language types and constraints on sequential learning and processing, suggesting that frequent language types tend to be those that have successfully adapted to these learning and processing limitations.

The final line of evidence supporting our explanation of basic word-order universals comes from a recent ALL study by our team. In one experiment, Christiansen took two of the grammars that Christiansen and Devlin had used for their network simulations—a consistent and an inconsistent grammar (see Table 16.1)—and trained forty subjects on sentences (represented as consonant strings) derived from the two grammars. Training and test

TABLE 16.1 *The two grammars used for stimuli generation in Christiansen (2000)*

Consistent Grammar		Inconsistent Grammar	
S	→ NP VP	S	→ NP VP
NP	→ (PP) N	NP	→ (PP) N
PP	→ NP post	PP	→ pre NP
VP	→ (PP) (NP) V	VP	→ (PP) (NP) V
NP	→ (PossP) N	NP	→ (PossP) N
PossP	→ NP Poss	PossP	→ Poss NP

Note: Vocabulary: {X, Z, Q, V, S, M}

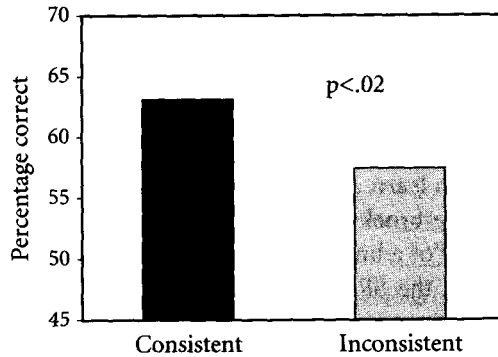


FIG. 16.5 The overall classification performance for the subjects trained on the consistent and inconsistent languages in Christiansen (2000).

materials were controlled for length and simple distributional differences. In the training phase of the experiment, subjects read and reproduced consonant strings on a computer. As in Christiansen *et al.* (in preparation), the subjects were not informed about the rule-based nature of the training items until they were about to commence the test phase.

The results are shown in Fig. 16.5. The twenty subjects trained on strings from the consistent grammar were significantly better at distinguishing grammatical from ungrammatical items than the twenty subjects trained on the inconsistent grammar. Together, Christiansen's ALL experiment and the three lines of evidence from Christiansen and Devlin converge to support our claim that basic word-order universals (head-ordering) can be explained in terms of non-linguistic constraints on sequential learning and processing. This research thus suggests that universal word-order correlations may emerge from non-linguistic constraints on learning, rather than being a product of innate linguistic knowledge. In the next section we show how constraints on complex question formation may be explained in a similar manner.

16.5 Subjacency without Universal Grammar

According to Pinker and Bloom (1990), subjacency is one of the classic examples of an arbitrary linguistic universal that makes sense only from

a linguistic perspective. Subjacency provides constraints on complex question formation. Informally, 'Subjacency, in effect, keeps rules from relating elements that are "too far apart from each other", where the distance apart is defined in terms of the number of designated nodes that there are between them' (Newmeyer 1991: 12). Consider the following sentences:

- (1) Sara heard (the) news that everybody likes cats.
 N V N comp N V N
- (2) What (did) Sara hear that everybody likes?
 Wh N V comp N V
- (3) *What (did) Sara hear (the) news that everybody likes?
 Wh N V N comp N V

According to the subjacency principle, sentence (3) is ungrammatical because too many boundary nodes are placed between the noun phrase complement and its respective 'gap'.

The subjacency principle, in effect, places certain restrictions on the ordering of words in complex questions. The movement of wh-items (*what* in Fig. 16.6) is limited with respect to the number of bounding nodes that it may cross during its upward movement. In English, the bounding nodes are S and NP (circled in Fig. 16.6). Put informally, as a wh-item moves up the




FIG. 16.6 A syntactic tree showing grammatical Wh-movement as in sentence 2.

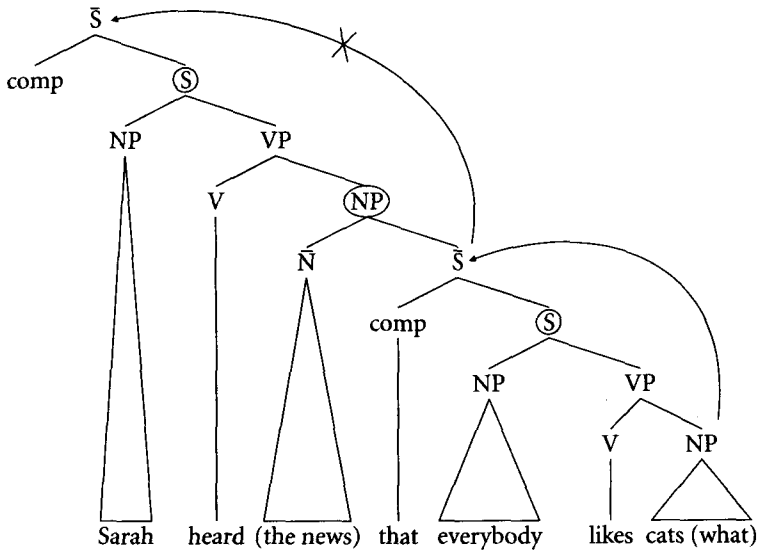


Fig. 16.7 A syntactic tree showing ungrammatical Wh-movement as in sentence 3.

tree it can use comps as temporary ‘landing sites’ from which to launch the next move. The subadjacency principle states that during any move only a single bounding node may be crossed. Sentence (2) is therefore grammatical because only one bounding node is crossed for each of the two moves to the top comp node. Sentence (3) is ungrammatical, however, because the wh-item has to cross two bounding nodes—NP and S—between the temporary comp landing site and the topmost comp, as illustrated in Fig. 16.7.

Not only do subadjacency violations occur in NP-complements, but they can also occur in wh-phrase complements. Consider the following examples:

- (4) Sara asked why everyone likes cats.
 N V Wh N V N
- (5) Who (did) Sara ask why everyone likes cats?
 Wh N V Wh N V N
- (6) *What (did) Sara ask why everyone likes?
 Wh N V Wh N V

According to the subadjacency principle, sentence 6 is ungrammatical because the interrogative pronoun has moved across too many bounding nodes (as was the case in sentence 3).

Ellefson and Christiansen (2000) explored an alternative explanation, suggesting that subadjacency violations are avoided, not because of a biological adaptation incorporating the subadjacency principle, but because language itself has undergone adaptations to root out such violations in response to non-linguistic constraints on sequential learning. We created two artificial languages to test this idea (including sentence types 1–6 above). As shown in Table 16.2, both languages consisted of six sentence types of which four were identical across the two languages. The two remaining sentence types involved complex question formation. In the natural language the two complex questions were formed in accordance with subadjacency, whereas the two complex questions in the unnatural language violated the subadjacency constraints. All training and test items were controlled for length and distributional information. As in the previous two ALL experiments, the twenty subjects trained in each condition were not told about the linguistic nature of the stimuli until they received the instructions for the test phase.

The results showed that the subjects trained on the natural grammar were significantly better at distinguishing grammatical from ungrammatical items than were the subjects trained on the unnatural language. As illustrated in Fig. 16.8, subjects in the natural condition were marginally better than the subjects in the unnatural condition at classifying strings related to the two complex questions. Interestingly, the natural group was significantly better at classifying the remaining four sentence types in comparison with the unnatural group—despite the fact that both groups were trained and tested on exactly the same general items. This suggests that the presence of the two unnatural question formation sentence types negatively affected the

TABLE 16.2 *The structure of the natural and unnatural languages in Ellefson and Christiansen (2000)*

Natural	Unnatural
N V N	N V N
Wh N V	Wh N V
N V N comp N V N	N V N comp N V N
N V Wh N V N	N V Wh N V N
Wh N V comp N V	*Wh N V N comp N V
Wh N V Wh N V N	*Wh N V Wh N V

Note: Vocabulary: {X, Z, Q, V, S, M}

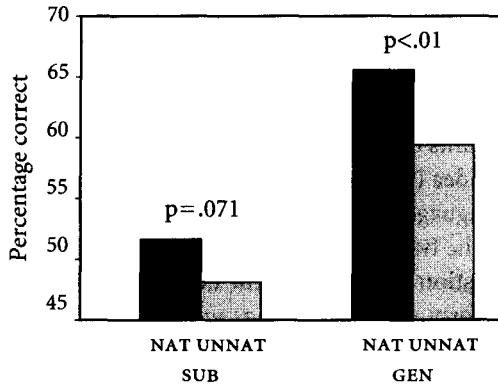


FIG. 16.8 The classification performance on subadjacency (SUB) and general (GEN) items for the subjects trained on the natural (NAT) and unnatural (UNNAT) languages in Ellefson and Christiansen (2000).

learning of the other four sentence types. In other words, the presence of the subadjacency violations in two of the sentence types in the unnatural language appears to have affected the learning of the language as a whole, not just the two complex question items. From the viewpoint of language evolution, languages such as this unnatural language would be likely to lose out in competition with other languages such as the natural language because the latter is easier to learn.

In principle, one could object that the reason why Ellefson and Christiansen found differences between the natural and the unnatural groups is because the former was in some way able to tap into an innately specified subadjacency principle when learning the language. Another possible objection is that the natural language follows the general pattern of English whereas the unnatural language does not, and that our human results could potentially reflect an 'English effect'. To counter these possible objections, and to support the suggestion that the difference in learnability between the two languages is brought about by constraints arising from sequential learning, Ellefson and Christiansen conducted a set of connectionist simulations of the human data using SRNs—a sequential learning device that clearly does not have subadjacency constraints built in. We used one network for each subject, and found that the networks were significantly better at learning the natural language than the unnatural language, as measured in terms of the ability to predict the correct sequence of elements in a string. Thus, the

simulation results closely mimicked the behavioural results, corroborating the suggestion that constraints on the learning and processing of sequential structure may explain why subadjacency violations tend to be avoided: these violations have been weeded out because they made the sequential structure of language too difficult to learn. Even though Ellefson and Christiansen's results do not capture all there is to subadjacency, they are nevertheless very encouraging, with future work expected to deal with other variations on subadjacency. Based on the current results we therefore venture to suggest that instead of an innate UG principle ruling out subadjacency violations, they may have been eliminated through linguistic adaptation.⁶

16.6 Conclusion

In this chapter we have argued that many of the universal constraints on the acquisition and processing of current languages are not linguistic in nature, but rather derive from underlying innate limitations on the learning and processing of hierarchically organized sequential structure. These cognitive constraints defined an important part of the niche within which languages have evolved through the adaptation of linguistic structure. In support of this perspective on language evolution we discussed evidence from three ALL studies. The first study, Christiansen *et al.* (in preparation), demonstrated that language breakdown in agrammatic aphasia is associated with impairment of sequential learning. Along with the other aphasia and neuroimaging studies we reviewed, this helps establish the direct link between language and sequential learning predicted by our account. The next study, Christiansen (2000), showed how constraints on sequential learning can explain basic word-order constraints. The third study, Ellefson and Christiansen (2000), provided a first step towards an explanation, based on sequential learning constraints, for why subadjacency violations tend to be avoided across the languages of the world. Together, the results from the three studies (and additional connectionist simulations) suggest that constraints arising from general cognitive processes, such as sequential learning and processing, are likely to play a larger role in sentence processing than

⁶ Note that whereas Berwick and Weinberg (1984) explain subadjacency as a consequence of processing constraints within a linguistically motivated parser, we provide an evolutionary explanation couched in terms of linguistic adaptation constrained to a large degree by *non-linguistic* limitations on sequential learning and processing.

has traditionally been assumed. What we observe today as linguistic universals may be stable states that have emerged through an extended process of linguistic evolution.

When language itself is viewed as a dynamic system sensitive to adaptive pressures, natural selection will favour combinations of linguistic constructions that can be acquired relatively easily given existing learning and processing mechanisms. Consequently, difficult to learn language fragments—such as Christiansen's (2000) inconsistent language and Ellefson and Christiansen's unnatural language—will tend to disappear. If we furthermore assume that the language production system is based conservatively on a processing system acquired in the service of comprehension, then this system would be unlikely to produce inconsistent grammatical structures or subadjacency violations because they would not be represented there in the first place. Thus, rather than having innate UG principles to ensure head-direction consistency or to rule out subadjacency violations, we argue that such linguistic universals derive from an evolutionary process of linguistic adaptation constrained by prior cognitive limitations on sequential learning and processing.

If language evolution is characterized primarily in terms of the adaptation of linguistic structure to cognitive constraints, it becomes imperative to determine the aspect(s) of language upon which natural selection works. One possibility is that selection takes place at the level of individual utterances; that is, only utterances with high fitness survive. For example, in an exploration of the trade-off between pressures from acquisition and production, Kirby (2001) shows that only utterances that have either a frequent meaning or a compositional syntax survive transmission from one generation to the next. Another possibility is that selection works at the level of grammars; that is, only grammars with high fitness survive. Briscoe (2000) presents simulations in which language evolution is couched in terms of changes over time in the distribution of different grammars that a population of learners acquires through exposure to a body of utterances produced by the previous generation. Based on the studies reported above, we propose that linguistic adaptation may be construed most fruitfully as involving a combination of utterance and whole-language selection. Properties of individual syntactic expressions, such as the degree of recursive inconsistency (as described above) or the frequency of occurrence (as in Kirby 2001), are likely to affect the process of linguistic adaptation. However, the evidence from Ellefson and Christiansen (2000) shows that the existence of less fit

utterance types in a language (e.g. the subadjacency items in the unnatural language), and not merely the unfit expressions themselves, can affect the learning of the language as whole. From this perspective, a language is more than just a collection of utterances. The relationship between utterances provides an additional source of variation upon which selection pressures can work in the adaptation of linguistic structure. Exactly how the single-utterance and whole-language selection pressures may interact is a question ripe for future research.

Finally, from a methodological perspective, it seems clear that ALL is a useful tool for exploring issues relating to language evolution. It may be objected that the languages used in ALL experiments are simple and deviate significantly from natural language. However, the same objection can be raised against the computational models of language evolution, but this has not diminished their impact, nor their usefulness to the study of language evolution. Moreover, ALL also provides a new tool with which to study other aspects of language evolution, such as creolization (Hudson and Newport 1998) and cross-species comparative aspects of language evolution (Hauser *et al.* 2001). In this way, ALL promises to open up a whole new direction in the search for evidence to rein in scientific theories of language evolution as well as to provide evidence for the adaptation of linguistic structure without linguistic constraints.

FURTHER READING

- CANGELOSI, A., and PARISI, D. (eds.) (forthcoming), *Simulating the Evolution of Language* (London: Springer Verlag). (Provides a series of chapters defining the state of the art in simulating the evolution of language.)
- GOMEZ, R. L., and GERKEN, L. A. (2000), 'Infant Artificial Language Learning and Language Acquisition', *Trends in Cognitive Sciences*, 4: 178–86. (Provides a cogent overview over the use of artificial language learning to study language acquisition.)
- MCMAHON, A. M. S. (1994), *Understanding Language Change* (Cambridge: Cambridge University Press). (Reviews the area of language change, including discussions of different notions of language as an organism.)
- MORGAN, J. L., MEIER, R. P., and NEWPORT, E. L. (1987), 'Structural Packaging in the Input to Language Learning: Contributions of Prosodic and Morphological Marking of Phrases to the Acquisition of Language', *Cognitive Psychology*, 19: 498–550. (Presents a series of artificial language learning experiments, demonstrating circumstances under which the presence of multiple linguistic cues, e.g. concord morphology, may help overcome some sequential learning difficulties.)

REFERENCES

- BATALI, J. (1998), 'Computational Simulations of the Emergence of Grammar', in J. R. Hurford, M. Studdert-Kennedy, and C. Knight (eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases* (Cambridge: Cambridge University Press), 405–26.
- BERWICK, R. C., and WEINBERG, A. S. (1984), *The Grammatical Basis of Linguistic Performance: Language Use and Acquisition* (Cambridge, Mass.: MIT Press).
- BICKERTON, D. (1995), *Language and Human Behavior* (Seattle: Washington University Press).
- BRISCOE, T. (2000), 'Macro and Micro Models of Linguistic Evolution', in J.-L. Dessalles and L. Ghadakpour (eds.), *Proceedings of the 3rd International Conference on the Evolution of Language* (Paris: École Nationale Supérieure des Télécommunications), 27–9.
- CHOMSKY, N. (1986), *Knowledge of Language* (New York: Praeger).
- CHRISTIANSEN, M. H. (1994), 'Infinite Languages, Finite Minds: Connectionism, Learning and Linguistic Structure', unpublished doctoral dissertation, Centre for Cognitive Science, University of Edinburgh.
- (2000), 'Using Artificial Language Learning to Study Language Evolution: Exploring the Emergence of Word Order Universals', in J. L. Dessalles and L. Ghadakpour (eds.), *Proceedings of the 3rd International Conference on the Evolution of Language* (Paris: École Nationale Supérieure des Télécommunications), 45–8.
- and DEVLIN, J. T. (1997), 'Recursive Inconsistencies are Hard to Learn: A Connectionist Perspective on Universal Word Order Correlations', in G. W. Cottrell (ed.) *Proceedings of the 19th Annual Cognitive Science Society Conference* (Mahwah, NJ: Lawrence Erlbaum), 113–18.
- KELLY, L., SHILLCOCK, R. C., and GREENFIELD, K. (in preparation). *Artificial Grammar Learning in Agrammatism*.
- CLEEREMANS, A. (1993), *Mechanisms of Implicit Learning: Connectionist Models of Sequence Processing* (Cambridge, Mass.: MIT Press).
- CORBALLIS, M. C. (1992), 'On the Evolution of Language and Generativity', *Cognition*, 44: 197–226.
- DAVIES, A. M. (1987), "'Organic" and "Organism" in Franz Bopp', in H. M. Hoenigswald and L. F. Wiener (eds.), *Biological Metaphor and Cladistic Classification* (Philadelphia: University of Pennsylvania Press), 81–107.
- DEACON, T. W. (1997), *The Symbolic Species: The Co-evolution of Language and the Brain* (New York: W. W. Norton).
- DOMINEY, P. F., HOEN, M., GOLEMBIOWSKI, M., and GUYOT, E. (2001), *Non-linguistic Training Transfers to Syntactic Comprehension Performance* (Tech. Memo No. 2001–2 (Lyons: Institut des Sciences Cognitives).
- DRYER, M. S. (1992), 'The Greenbergian Word Order Correlations', *Language*, 68: 81–138.

- ELLEFSON, M. R., and CHRISTIANSEN, M. H. (2000), 'Subjacency Constraints Without Universal Grammar: Evidence from Artificial Language Learning and Connectionist Modeling', in L. R. Gleitman and A. K. Joshi (eds.), *The Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (Mahwah, NJ: Erlbaum), 645–50.
- ELMAN, J. L. (1990), 'Finding Structure in Time', *Cognitive Science*, 14: 179–211.
- GIVÓN, T. (1998), 'On the Co-evolution of Language, Mind and Brain', *Evolution of Communication*, 2: 45–116.
- GOMEZ, R. L., and GERKEN, L. A. (2000), 'Infant Artificial Language Learning and Language Acquisition', *Trends in Cognitive Sciences*, 4: 178–86.
- GROSSMAN, M. (1980), 'A Central Processor for Hierarchically Structured Material: Evidence from Broca's Aphasia', *Neuropsychologia*, 18: 299–308.
- HAUSER, M. D., NEWPORT, E. L., and ASLIN, R. N. (2001), 'Segmentation of the Speech Stream in a Non-human Primate: Statistical Learning in Cotton-Top Tamarins', *Cognition*, 78: B53–B64.
- HAWKINS, J. A. (1994), *A Performance Theory of Order and Constituency* (Cambridge: Cambridge University Press).
- HUDSON, C. L., and NEWPORT, E. L. (1998), 'Creolization: Could Adults Really Have Done it All?', in A. Greenhill, H. Littlefield, and C. Tano (eds.), *Proceedings of the 23rd Annual Boston University Conference on Language Development* (Boston: Cascadilla Press), 265–76.
- KIMURA, D. (1988), 'Review of What the Hands Reveal about the Brain', *Language and Speech*, 31: 375–8.
- KIRBY, S. (1998), 'Fitness and the Selective Adaptation of Language', in J. R. Hurford, M. Studdert-Kennedy, and C. Knight (eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases* (Cambridge: Cambridge University Press), 359–83.
- (2000), 'Language Evolution Without Natural Selection: From Vocabulary to Syntax in a Population of Learners', in C. Knight, J. R. Hurford, and M. Studdert-Kennedy (eds.), *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form* (Cambridge: Cambridge University Press), 303–33.
- (2001), 'Spontaneous Evolution of Linguistic Structure: An Iterated Learning Model of the Emergence of Regularity and Irregularity', *IEEE Transactions on Evolutionary Computation*, 5/2: 102–10.
- LAKOFF, G. (1987), *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind* (Chicago: University of Chicago Press).
- LANGACKER, R. W. (1987), *Foundations of Cognitive Grammar: Theoretical Perspectives* (Stanford, Calif.: Stanford University Press), i.
- MCMAHON, A. M. S. (1994), *Understanding Language Change* (Cambridge: Cambridge University Press).
- MAESS, B., KOELSCH, S., GUNTER, T. C., and FRIEDERICI, A. D. (2001), 'Musical Syntax Is Processed in Broca's Area: An MEG study', *Nature Neuroscience*, 4: 540–5.

- MORGAN, J. L., MEIER, R. P., and NEWPORT, E. L. (1987), 'Structural Packaging in the Input to Language Learning: Contributions of Prosodic and Morphological Marking of Phrases to the Acquisition of Language', *Cognitive Psychology*, 19: 498-550.
- NERLICH, B. (1989), 'The Evolution of the Concept of 'Linguistic Evolution' in the 19th and 20th century', *Lingua*, 77: 101-12.
- NEWMAYER, F. (1991), 'Functional Explanation in Linguistics and the Origins of Language', *Language and Communication*, 11: 3-28.
- OSTERHOUT, L., and HOLCOMB, P. J. (1992), 'Event-Related Brain Potentials Elicited by Syntactic Anomaly', *Journal of Memory and Language*, 31: 785-806.
- PATEL, A. D., GIBSON, E., RATNER, J., BESSON, M., and HOLCOMB, P. J. (1998), 'Processing Syntactic Relations in Language and Music: an Event-Related Potential Study', *Journal of Cognitive Neuroscience*, 10: 717-33.
- PERCIVAL, W. K. (1987), 'Biological Analogy in the Study of Languages before the Advent of Comparative Grammar', in H. M. Hoenigswald and L. F. Wiener (eds.), *Biological Metaphor and Cladistic Classification* (Philadelphia: University of Pennsylvania Press), 3-38.
- PIATELLI-PALMARINI, M. (1989), 'Evolution, Selection and Cognition: From "Learning" to Parameter Setting in Biology and in the Study of Language', *Cognition*, 31: 1-44.
- PINKER, S. (1994), *The Language Instinct: How the Mind Creates Language* (New York: William Morrow).
- and BLOOM, P. (1990), 'Natural Language and Natural Selection', *Brain and Behavioral Sciences*, 13: 707-27.
- SCHLEICHER, A. (1863), *Die Darwinsche Theorie Und Die Sprachwissenschaft* (Weimar: Böhlau).
- STEINHAEUER, K., FRIEDERICI, A. D., and PFEIFER, E. (2001), *ERP Recordings While Listening to Syntax Errors in an Artificial Language: Evidence from Trained and Untrained Subjects*, Poster presented at the 14th Annual CUNY Conference on Human Sentence Processing, Philadelphia.
- STEVICK, R. D. (1963), 'The Biological Model and Historical Linguistics', *Language*, 39: 159-69.
- VAN EVERBROECK, E. (1999), 'Language Type Frequency and Learnability: A Connectionist Appraisal', in M. Hahn and S. C. Stoness (eds.), *Proceedings of the 21st Annual Cognitive Science Society Conference* (Mahwah, NJ: Erlbaum), 755-60.