Language Evolution and Change

Morten H. Christiansen and Rick Dale

Introduction

No direct evidence remains from before the emergence of writing systems to inform theories about the evolution of language. Only as evidence is amassed 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. To go beyond existing data, rigorously controlled thought experiments can be used as crucial tests of competing theories. Computational modeling has become a valuable resource for such tests because it enables researchers to test hypotheses about specific aspects of language evolution under controlled circumstances (Cangelosi and Parisi, 2002; Turner, 2002). With the help of computational simulations, it is possible to study various processes that may have been involved in the evolution of language, as well as the biological and cultural constraints that may have shaped language into its current form (see Evolution and Learning in Neural Networks).

Connectionist models have played an important role in the computational modeling of language evolution. In some cases, the networks are used as simulated agents to study how social transmission via learning might give rise to the evolution of structured communication systems. In other cases, the specific properties of neural network learning are enlisted to help illuminate the constraints and processes that may have been involved in the evolution of language. This article surveys this connectionist research, starting from the emergence of early syntax and continuing to the role of social interaction and constraints on network learning in subsequent evolution of language and to linguistic change within existing languages.

Emergence of Simple Syntax

Models of language evolution focus on two primary questions: how language emerged, and how languages continue to change over time. An important feature of the first question is the emergence of syntactic communication. Cangelosi (1999) studied the evolution of simple communication systems, but with an emphasis on the emergence of associations not only between objects (meaning) and symbols (signal), but also between the symbols themselves (syntax). In particular, the aim was to demonstrate that simple syntactic relations (a verb-object rule) could evolve through a combination of communicative interactions and cross-generational learning in populations of neural networks.

In Cangelosi's simulations, populations of networks evolved based on their ability to forage in an environment consisting of a two-dimensional 100×100 array of cells. About 12% of the cells contained randomly placed mushrooms that served as food. Three types of mushrooms were edible, increasing a network's fitness if collected, whereas another three types were poisonous, decreasing the network's fitness if collected. The networks had a standard feed–forward architecture with a single hidden unit layer and were trained using backpropagation (see Backpropagation: General Principles and Issues for Biology). Input was represented in terms of three sets of input units encoding the location of a mushroom, the visual features of the mushroom, and words naming objects or actions. The output contained sets of units representing actions (*approach, avoid, discriminate*) and words with the latter units organized into two winner-take-all clusters (object and verb). Populations consisted of 80 networks, each with a life span of 1,000 actions. The 20 networks with the highest fitness level were selected for asexual reproduction, each producing four offspring through random mutation of 10% of its starting weights. During the first 300 generations, the populations evolved an ability to discriminate between edible and poisonous mushrooms without the use of words. In subsequent populations, parents provided teaching input for the learning of words denoting the different mushrooms (objects) and the proper action to take (verbs). The simulations were repeated with different random starting populations. Sixty-one percent of the simulations resulted in optimal vocabulary acquisition, with different "verb" symbols used with edible (*approach*) and poisonous (*avoid*) mushrooms, and different "noun" symbols used for the different types of mushrooms.

The simulations indicate how a simple noun-verb communication system can evolve in a population of networks. Because the features of a mushroom were perceived only 10% of the time, paying attention to the parental language input provided a selective advantage with respect to foraging, thus reinforcing successful linguistic performance.

Another approach to the emergence of elementary syntax has been offered by Batali (1998). He suggested that a process of negotiation between agents in a social group may have given rise to coordinated communication. Whereas Cangelosi's model involved the emergence of rudimentary verb-object syntax in a foraging environment, Batali's networks were assigned the task of mapping meaning onto a sequence of characters for the purpose of communication in a social environment. The networks in this simulation did not start out with a predetermined syntactic system. Instead, a process of negotiation across generations engendered the evolution of a syntactic system to convey common meanings.

Each agent in the simulation was a simple recurrent network (SRN; Elman, 1990), capable of processing input sequences consisting of four characters and producing an output vector representing a meaning involving a subject and a predicate. In a negotiation round, one network was chosen as a learner, and ten randomly selected teachers conveyed a meaning converted into a string of characters. The learner then processed the string produced by the teacher, and was trained using the difference between the teacher's and the learner's meaning vectors. Batali described this interaction between learners and teachers as a kind of negotiation, since each must adjust weights in accordance with its own cognitive state and that of others. At the start of the simulations the networks generated only very long strings that were unique to each meaning. After several thousand rounds of negotiation, the agents developed a more efficient and partially compositional communication system, with short sequences of letters used for particular predicates and referents. To test whether novel meanings could be encoded by the communication system, Batali omitted ten meanings, and reran the simulations. After training, networks performed well at sending and processing the omitted meaning vectors, demonstrating that the rudimentary grammar exhibited systematicity that accommodated a structured semantics.

Batali's model offers illuminating observations for the evolution of language. An assumption of this model was that social animals can use their own cognitive responses (in this case, translating meaning vectors into communicable signals) to predict the cognitive state of other members of their community. Batali compared this ability to one that may have arisen early in hominids and contributed to the emergence of systematic communication. Once such an elementary communication system is in place, migration patterns may have promoted dialectical variations. The next section explores how linguistic diversity might arise as a result of geographical separation between groups of communicating agents.

Linguistic Diversity

The diversity of the world's many languages has offered puzzling questions for centuries. Computational simulations allow for the investigation of factors influencing the distribution and diversity of language types. An intuitive approach, considered in this section, is that languages assume an adaptive shape governed by various constraints in the organism and environment. Livingstone and Fyfe (1999) have proposed an alternative perspective based on simulations in which linguistic diversity arises simply as a consequence of spatial organization and imperfect language transmission in a social group.

The social group in the simulation consisted of networks with two layers of three input and output units, bidirectionally connected and randomly initialized. As in Batali's simulations, agents were given the task of mapping a meaning vector onto an external "linguistic" signal. For each generation, a learner and a teacher were randomly selected. The output of the teacher was presented to the learner, and the error between meaning vectors was used to change the learner's weights. Each successive generation had agents from the previous generation acting as teachers. The agents were spatially organized along a single dimension and communicated only with other agents within a fixed distance. By comparing agents across this spatial organization, performance akin to a dialect continuum was observed: small clusters of agents communicated readily, but as the distance among them increased, error in communication increased. When the simulation was implemented without spatial organization (i.e., each agent was equally likely to communicate with all others), the entire population quickly negotiated a global language, and diversity was lost. This model supports the position that diversity is a consequence of spatial organization and imperfect cultural transmission.

The results of Livingstone and Fyfe's as well as Batali's simulations may not rely directly on the properties of neural network learning, but rather on the processes of learning-based social transmission. However, when it comes to explaining why certain linguistic forms have become more frequent than others, the specific constraints on learning in such networks come to the fore. The next section discusses how limitations on network learning can help explain the existence of certain so-called linguistic universals.

Learning-Based Linguistic Universals

Despite the considerable diversity that can be observed across the languages of the world, it is also clear that languages share a number of relatively invariant features in the way words are put together to form sentences. Spatial organization and error in transmission cannot account for these widespread commonalities. Instead, the specific constraints on neural network learning may offer explanations for these consistent patterns in language types. As an example, we can consider heads of phrases, that is, the particular word in a phrase that determines the properties and meaning of the phrase as a whole (such as the noun *boy* in the noun-phrase *the boy with the bicycle*). Across the world's languages, there is a statistical tendency toward a basic format in which the head of a phrase consistently is 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 noun-phrase in a transitive verb phrase such as *eat curry*. In contrast, speakers of Hindi would say the equivalent of *curry eat*, because Hindi is a head-last language.

Christiansen and Devlin (1997) trained SRNs with eight input and eight output units encoding basic lexical categories (i.e., nouns, verbs, prepositions, and a possessive genitive marker) on corpora generated by 32 different grammars with differing amount of head-order consistency. The networks were trained to predict the next lexical category in a sentence. Importantly, these networks did not have built-in linguistic biases; rather, they were biased toward the learning of complex sequential structure. 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 in a given grammar and the degree to which the network had learned to master the grammatical regularities underlying that grammar. The higher the inconsistency, the more erroneous the final network performance was. The sequential biases of the networks made the corpora generated by consistent grammars considerably easier to acquire than the corpora generated by inconsistent grammars. Christiansen and Devlin further collected frequency data concerning the specific syntactical constructions used in the simulations. They found that languages incorporating fragments that the networks found hard to learn tended to be less frequent than languages the network learned more easily. This suggests that constraints on basic word order may derive from nonlinguistic constraints on the learning and processing of complex sequential structure. Grammatical constructions incorporating a high degree of head-order inconsistency may simply be too hard to learn, and would therefore tend to disappear.

More recently, Van Everbroeck (1999) presented network simulations in a similar vein in support of an explanation for language-type frequencies based on processing constraints. He trained recurrent networks (a variation on the SRN) to produce the correct grammatical role assignments for noun-verb-noun sentences that were presented one word at a time. The networks had 26 input units, providing distributed representations of nouns and verbs as well as encodings of case markers, and 48 output units, encoding the distributed noun-verb representation according to grammatical role. Forty-two different language types were used to represent cross-linguistic variation in three dimensions: word order (e.g., subject-verb-object), and noun and verb inflection. The results of the simulations coincided with many observed trends in the distribution of the world's languages. Subject-first languages, both of which make up the majority of language types (51% and 23%, respectively), were easily processed by the networks. Object-first languages, on the other hand, were not well processed, and they have very low frequency in the world's languages (object-verb-subject: 0.75%; object-subject-verb: 0.25%). Van Everbroeck argued that these results were a predictable product of network processing constraints. Not all results, however, were directly proportional to actual language-type frequencies. For example, verb-subject-object languages account for only 10% of the world's language types, but the model's performance on it exceeded performance on the more frequent subject-first languages. Van Everbroeck suggested that making the simulations more sophisticated (incorporating semantics or other aspects of language) might allow network performance to better approach observed frequencies. Together, the simulations by Van Everbroeck and by Christiansen and Devlin provide preliminary support for a connection between learnability and frequency in the world's languages based on the learning and processing properties of connectionist networks. The next section discusses additional simulations that show how similar network properties may also help explain linguistic change within a particular language.

Linguistic Change

The English system of verb inflection has changed considerably over the past 1,100 years. Simulations by Hare and Elman (1995) demonstrate how neural network learning and processing constraints may help explain the observed pattern of change. The morphological system of Old English (ca. 870) was quite complex, involving at least ten different classes of verb inflection (with a minimum of six of these being "strong"). The simulations involved several "generations" of neural networks, each of which received as input the output generated by a trained net from the previous generation. The first net was trained on data representative of the verb classes from Old English. However, training was stopped before learning could reach optimal performance. This reflected the causal role of imperfect transmission in language change. The imperfect output of the first net was used as input for a second generation net, for which training was also halted before learning reached asymptote. Output from the second net was then given as input to a third net, and so on, until seven generations were trained. This training regime led to a gradual change in the morphological system. These changes can be explained by verb frequency in the training corpus, and internal phonological consistency (i.e., distance in phonological space between prototypes). The results revealed that membership in small classes, inconsistent phonological characteristics, and low frequency all contributed to rapid morphological change. As the morphological system changed through generations in these simulations, the pattern of results closely resembled the historical change in English verb inflection from a complex past tense system to a dominant "regular" class and small classes of "irregular" verbs.

Discussion

This article has surveyed the use of neural networks for the modeling of language evolution and change. The results discussed here are encouraging, even though neural network modeling of language evolution is very much in its infancy. However, it is also clear that the current models suffer from obvious shortcomings. Most of them are highly simple and do not fully capture the vast complexity of the issues at hand. For example, the models of the emergence of verb-object syntax and linguistic diversity incorporated very simple relationships between meaning and form. Moreover, although the simulations of the influence of processing constraints on the shape of language involved relatively complex grammars, they did not include any relationship between the language system and the world. Nevertheless, these models demonstrate the potential for exploring the evolution of language from a computational perspective.

Both connectionist and nonconnectionist models (e.g., Nowak and Komarova, 2001) have been used to provide important thought experiments in support of theories of language evolution. Connectionist models have become prominent in such modeling, both for their ability to simulate social interaction in populations and for their demonstrations of how learning constraints imposed on communication systems can engender many of the linguistic properties we observe today. Together, the models point to an important role for cultural transmission in the origin and evolution of language. This perspective receives further support from neuroscientific considerations, suggesting a picture of language and brain that argues for their co-evolution (e.g., Deacon, 1997). The studies discussed here highlight the promise of neural network approaches to these issues. Future studies will likely seek to overcome current shortcomings and move toward more sophisticated simulations of the origin and evolution of language.

Roadmap: Linguistics and Speech Processing; Neuroethology and Evolution
Background: Language Processing
Related Reading: Constituency & Recursion in Language; Evolution and Learning in Neural Networks; Language Evolution, The Mirror System Hypothesis

References

Batali, J., 1998, Computational simulations of the emergence of grammar, in *Approaches to the Evolution of Language: Social and Cognitive Bases* (J. R. Hurford, M. Studdert-Kennedy, and C. Knight, Eds.), Cambridge, Engl.: Cambridge University Press, pp. 405–426.

Cangelosi, A., 1999, Modeling the evolution of communication: From stimulus associations to grounded symbolic associations, in *Advances in Artificial Life: Proceedings of the ECAL99 European Conference on Artificial Life* (D. Floreano, J. Nicoud, and F. Mondada, Eds.), Berlin: Springer-Verlag, pp. 654–663.

Cangelosi, A., and Parisi, D., 2002, Computer simulation: A new scientific approach to study of language evolution, in *Simulating Language Evolution* (A. Cangelosi and D. Parisi, Eds.), London: Springer-Verlag, pp. 3–28. ◆

Christiansen, M. H., and Devlin, J. T., 1997, Recursive inconsistencies are hard to learn: A connectionist perspective on universal word order correlations, in *Proceedings of the 19th Annual Cognitive Science Society Conference*, Mahwah, NJ: Erlbaum, pp. 113–118.

Deacon, T., 1997, *The Symbolic Species: The Co-evolution of Language and the Brain*, New York: Norton. ◆

Elman, J. L., 1990, Finding structure in time, *Cogn. Sci.*, 14:179–211.

Hare, M., and Elman, J. L., 1995, Learning and morphological change, *Cognition*, 56:61–98.

Nowak, M. A., and Komarova, N. L., 2001, Towards an evolutionary theory of language, *Trends Cognitive Sci.*, 5:288–295.

Livingstone, D., and Fyfe, C., 1999, Modelling the evolution of linguistic diversity, in *Advances in Artificial Life: Proceedings of free ECAL99 European Conference on Artificial Life* (D. Floreano, J. Nicoud, and F. Mondada, Eds.), Berlin: Springer–Verlag, pp. 704–708.

Turner, H., 2002, An introduction to methods for simulating the evolution of language, in *Simulating Language Evolution* (A. Cangelosi and D. Parisi, Eds.), London: Springer-Verlag, pp. 29–50. ◆

Van Everbroeck, E., 1999, Language type frequency and learnability: A connectionist appraisal, in *Proceedings of the 21st Annual Cognitive Science Society Conference*, Mahwah, NJ: Erlbaum, pp. 755–760.