

Book Reviews

Linguistic Evolution through Language Acquisition: Formal and Computational Models

Ted Briscoe (editor)
(University of Cambridge)

Cambridge, England: Cambridge
University Press, 2002, vii+349 pp;
hardbound, ISBN 0-521-66299-0, \$65.00

Reviewed by
Michael A. Arbib
University of Southern California

This book is really two books, which do not communicate with each other once one gets past the editor's introduction in chapter 1. The chapters by Worden, Batali, Kirby, and Hurford (see Table 1), whom I shall refer to collectively as WBKH, use computer simulation to demonstrate that agents with no innate syntactic structure can interact to create and preserve both the lexicon and syntax of languages over many generations. On the other hand, the chapters by Niyogi, Turkel, and Briscoe (collectively, NTB) all base their simulations on an innate Universal Grammar—the general form of grammar is built into the agents, and the lexicon is given little or no importance. I shall briefly outline the other two chapters, by Oliphant and by Steels and Kaplan, in the context of the WBKH work below.

All the authors assume that the agents under study have a great deal of “innate” language-related structure. The “linguistic evolution” studied here is not the evolu-

Table 1

Authors and titles of chapters in *Linguistic Evolution through Language Acquisition*.

1. Ted Briscoe: Introduction
2. Michael Oliphant: Learned systems of arbitrary reference: The foundation of human linguistic uniqueness
3. Luc Steels and Frederic Kaplan: Bootstrapping grounded word semantics
4. Robert Worden: Linguistic structure and the evolution of words
5. John Batali: The negotiation and acquisition of recursive grammars as a result of competition among exemplars
6. Simon Kirby: Learning, bottlenecks and the evolution of recursive syntax
7. Partha Niyogi: Theories of cultural evolution and their application to language change
8. William J. Turkel: The learning guided evolution of natural language
9. Ted Briscoe: Grammatical acquisition and linguistic selection
10. James R. Hurford: Expression/induction models of language evolution: Dimensions and issues

tion of the *capacity* for language, but rather the ability of a community to build on innate mechanisms to seek coherence in the encoding of meanings by strings of symbols (WBHK) or the parameter settings in their grammars (NTB). The phrase *through language acquisition* in the title means that as each agent is added to a population, it seeks to model its sample of the utterances of the existing population, and in so doing changes the population profile. The language defined by the statistics of the population's output "evolves" accordingly.

For NTB, all study starts with an innate Universal Grammar characterized by a set of principles and parameters. The job of the language learner is to infer the parameter settings of the strings that it encounters, or to reach consensus within a population of agents as to which parameter settings to converge upon in resolving initially disparate data—data that take the form of strings of syntactic categories rather than words whose meaning or categorization is to be assigned or discovered.

For WBKH, each agent has a set M of structured meanings (e.g., semantic trees or logical expressions over some small set of symbols) and a set S of strings of symbols. (Oliphant addresses the trivial case of agents with a fixed, finite set of unstructured meanings, each of which is to be paired with a single symbol from another finite set.) Moreover, WBHK assume that the agents in successive generations have innate capacities that transcend the abilities of nonhuman creatures: In addition to the set S of strings and structured representations for the set M of shared meanings, they are also equipped with learning algorithms whose sole purpose is to infer rules that can generate meaning-to-string encodings that increasingly match the encodings used by other individuals. Further, WBKH assume that the learner has direct access to the meaning of each string; that is, its training data are pairs of the form $\langle \text{meaning}, \text{string} \rangle$. (Steels and Kaplan address the problem that in general the meaning of a string is not explicit, and so the hearer may mistakenly pair the currently heard string with the wrong meaning. Does a phrase uttered when looking at a red triangle refer to its size, shape, color, or location, or to some combination of these properties? Steels and Kaplan show how patterns of shared attention across many items may eventually resolve such ambiguities.) A population creates and learns meaning-string pairs. At first, the assignment of string to meaning is essentially random and varies from individual to individual, but as each new member is added to a population, it will seek to learn pairings that are already in use.

WBHK show that as "old" agents leave the population and "new" agents join it, something like a coherent language emerges across the course of generations. The resulting consensus assignment gains its power by reflecting semantic structure in the structure of the corresponding string. The key point is emergence of a pattern of coherence that has never existed before in the history of the "species." Compositionality and recursion in the mapping emerge across generations as a result of general conditions of learnability and coherence, rather than being built in as innate principles. The accounts given by Worden, Batali, and Kirby of the emergence of language structure are similar in their general spirit, but each contains different, and rather strong, assumptions about the structure that the agents are provided. Hurford provides a useful synthesis by contrasting the assumptions made by Batali and Kirby in this volume and those made in earlier papers by Batali, Kirby, and Hurford himself. I found the various WBHK mechanisms interesting—though I see much further work required to explain the biological evolution whereby hominids came to have mechanisms for mapping open sets of meanings to symbol strings, and I would like to see models exploring how language might guide the discovery of concepts, rather than expressing concepts that are already built in, let alone formalized.

The NTB approach is essentially a generalization of the following idea (which I think is not an unfair caricature of the theory offered by Niyogi): "Suppose a child has in its head a grammar for both English and German. Here is an algorithm whereby the child can decide whether a set of strings of non-terminals is better interpreted as having an English or German setting of parameters in a Universal Grammar." I confess that I cannot get excited about such a (to me) implausible view of language acquisition. Turkel offers an interesting analysis of the merits of evolving agents that have fixed certain parameters (principles) while leaving other parameters open to learning. Briscoe (in the least readable chapter of the book—the editor needed a good editor) offers a general Bayesian algorithm for inferring the most likely setting of n parameters in a Universal Grammar that is based on categorial grammar.

A strange feature of the book is that it is essentially data free. The WBKH methodology is this: "Here is a set of interacting agents. Here are some measures of how the 'languages' of the agents in a population do or do not cohere with one another. Here are simulation results showing interesting patterns of increasing coherence across the generations." The patterns are indeed interesting but there are no studies of this form: "Here is a model of language acquisition. Here are real data on language acquisition in human children. Look how well the model explains key aspects of the data. It is plausible that these mechanisms have long been present in hominids and so could account for historical patterns of the emergence of language as well as its ontogeny in the individual child." More encouragingly, NTB do offer hints of how one might use historical data on languages with studies of this form: "Here is a model of language acquisition. Here are real data on language change in an attested population across the generations. Look how well the model explains key aspects of these patterns of change." Briscoe uses his model of parameter setting to provide a computational framework for Bickerton's bioprogram hypothesis on the rapid emergence of Hawaiian creole. Since Bickerton's hypothesis is controversial, I look forward to further research by Briscoe that looks at data on a wide range of creoles and assesses the extent to which his modeling might tip the balance between the universalist and substratum hypotheses (Muysken and Smith 1986). Niyogi promises to consider the evolution of English from the 9th to 14th centuries A.D., but in the end gives a model of how to flip a couple of parameters rather than giving a generation-by-generation account of the cumulative changes whereby a new language emerged. This only strengthens my suspicion that characterizing a language by innate parameters is too coarse an instrument to describe historical change, and that the WBKH type of microanalysis, whereby small details can come and go, is far more promising for the study of linguistic evolution through language acquisition.

Reference

Muysken, Pieter C. and N. Norval Smith,
editors. 1986. *Substrata versus Universals in
Creole Genesis*. John Benjamins,
Amsterdam.

Michael A. Arbib is the Fletcher Jones Professor of Computer Science, as well as a Professor of Biological Sciences, Biomedical Engineering, Electrical Engineering, Neuroscience and Psychology, at the University of Southern California. His books related to language include *Neural Models of Language Processes* (edited with D. Caplan and J. C. Marshall, 1982), *From Schema Theory to Language* (with E. J. Conklin and J. C. Hill, 1987), *An Introduction to Formal Language Theory* (with A. J. Kfourri, R. N. Moll, and J. Pustejovsky, 1988), and *Beyond the Mirror: Biology and Culture*

in the Evolution of Brain and Language (Oxford University Press, to appear). He is also editor of *The Handbook of Brain Theory and Neural Networks* (second edition, The MIT Press, 2003), which contains a number of articles on language processing. Arbib's address is USC Brain Project, University of Southern California, Los Angeles, CA 90089-2520; e-mail: arbib@pollux.usc.edu.

Implementing Typed Feature Structure Grammars

Ann Copestake

(University of Cambridge)

Stanford, CA: CSLI Publications (CSLI lecture notes, number 110), 2002, xi+223 pp; distributed by the University of Chicago Press; hardbound, ISBN 1-57586-261-1, \$62.00; paperbound, ISBN 1-57586-260-3, \$22.00

Reviewed by

Gerald Penn

University of Toronto

Typed feature logic had been promising to deliver large-scale, semantically rich grammars ever since Pollard and Sag's (1987) seminal book on head-driven phrase structure grammar (HPSG) appeared. Several private research labs had been writing them, we were told, but they were so good and so valuable that no, sorry, you can't take a peek at them. It was not until the Linguistic Grammars Online (LinGO) consortium, of which Copestake is a member, dedicated themselves to the task of writing a publicly available English grammar, now known as the *English Resource Grammar* or *ERG* (Flickinger 1999), that typed feature logic finally made good on that promise. There is now at least one large-scale grammar based on typed feature structures.

The consortium extended an existing system for developing lexical knowledge bases, the LKB, to the task of developing such a grammar. The LKB is not the first grammar development system for typed feature logic, nor the fastest, nor the most expressive (all three of which are rumored to exist within the walls of the aforementioned research labs), but it, too, is publicly available. The intention of Copestake's book, about half of which is a user's manual for the LKB, is to make it even more accessible, particularly to those, as the author puts it, "who do not like equations, mathematical symbols, Greek letters and so on" (p. 4). That is indeed a tall order, and skeptics could certainly argue that this is not an audience from whom we would generally expect precise, large-scale grammars anyway. Copestake does a very respectable job of bringing the necessary discrete mathematics to them, however, mainly by illustrating concepts through many well-chosen examples and exercises throughout the presentation. To that extent, this book will serve as a very welcome introduction to the topic for novice grammar writers.

The book also contains a number of optional sections that present the definitions more formally, for those who wish to delve further into the logic itself. Even in these there are some ambiguities, though, as well as some mistakes and parochial choices of terminology. The distinction between *specifications* of mathematical structures and the structures themselves is often blurred, for example. The phrase *constraint on a type* is used sometimes to refer to a specific statement of entailment that a grammar writer provides (e.g., p. 68), but elsewhere to refer to the most general satisfier of a type relative to the entire constraint system, which is computed under multiple inheritance from those specifications (p. 75). *Well-formedness*, which is normally used to refer to the syntactic correctness of feature structures, is instead used here to refer to well-typing, which indicates an adherence to the semantic type system of the grammar. The

syntax of descriptions is presented in Backus-Naur form, but descriptions themselves are not abstractly defined (as distinct from the feature structures they denote). In an intensional description logic such as typed feature logic, this distinction is important, and one that beginning students often miss. These ambiguities, combined with a very sparse collection of bibliographic citations, may make it difficult for novices to find clarification elsewhere.

A greater concern is that, among the expositions of syntax and common pitfalls that constitute grammar training, there is very little to be found here in the way of grammar guidance. What do not just functioning but *well-written* grammars look like? What properties should that compliment—that a grammar is well written—imply? If some connection to the theory of programming languages can be drawn, a grammar that is well written is a grammar that is efficient, but it is also a grammar that is documented and testable or otherwise verifiable, a grammar that is reusable and adaptable, and to that end, a grammar expressed in terms that are appropriately modular and encapsulated. These topics are only briefly discussed in the penultimate chapter, “Advanced Features,” and even then only with an acknowledgment (p. 192) that they are indeed important, and that semantic typing helps.

Of course, only time will tell whether, from among the several generations of LKB users that will undoubtedly issue from this book, truly exceptional HPSG writers will emerge. Historically, at least, the proposals for capturing the functionality that large, well-written grammars require have encompassed a rather wide range of grammatical constructs and their combinations. The LKB, which essentially consists of only phrase structure rules and a large partial ordering of types, falls very close to one end of that spectrum. Although the observation that a grammar of the size and coverage of the English Resource Grammar should have first been developed at that end certainly has not been lost on those of us who have proposed alternatives, even a casual glance at the now complete ERG must cause one to speculate whether every grammar realized within the restricted functionality of the LKB is automatically a well-written one. More attention to principles of grammar implementation was therefore probably warranted here.

In simultaneously designing the English Resource Grammar and the LKB system, the LinGO consortium had the rare opportunity not only to experiment with other answers to this question, but to set a higher standard among grammar development systems in the functionality required to do a proper job of large-scale grammar design. My own recollection of the consortium’s reasoning, however, is that, while this design decision was motivated by the requirements of the English Resource Grammar, those requirements in turn were motivated by the intention of making the grammar as portable to other grammar development systems as possible. This resulted in a least upper bound (pun intended) of the functionality commonly available in these systems in the early 1990s, which in retrospect may have been an unfortunate choice. This is clearly an indictment of a much broader research program than what Copestake’s book represents, but to a great extent the book and the system it describes are very much reflections of both the positive and negative outcomes of that program. In that respect, the book is indeed a very faithful record and makes these outcomes very accessible for a wide audience to judge for themselves.

References

- Flickinger, Dan. 1999. The LinGO English Resource Grammar. Available at (<http://lingo.stanford.edu>).
- Pollard, Carl and Ivan Sag. 1987. *Information-Based Syntax and Semantics. Volume 1: Fundamentals*. Stanford, CA: Center for the Study of Language and Information.

Gerald Penn is the co-designer of the Attribute Logic Engine (ALE). He has conducted research in typed feature logics and HPSG for over 10 years. Penn's address is Department of Computer Science, University of Toronto, 10 King's College Road, Toronto M5S 3G4, Canada; e-mail: gpenn@cs.toronto.edu; URL: www.cs.toronto.edu/~gpenn.