A Computational Model of First Language Acquisition

Nobuo Satake
(University of Tokyo)

Singapore: World Scientific, 1990,
viii + 197 pp.
(World Scientific Series in Computer
Science 20)
Hardbound ISBN 981-02-0139-7, $28.00

Reviewed by
Robert C. Berwick
Massachusetts Institute of Technology

In a recent survey of early language acquisition, Gleitman, Gleitman, Landau, and Wanner (1988) cite Leonard Bloomfield (1933, p. 29) as remarking that "language learning is doubtless the greatest intellectual feat any one of us is ever required to perform." Given this, it is equally no small feat to attempt to build a computer model that does the same thing. Satake's book is one of but a handful of attempts in the computational linguistics tradition to take up this challenge—somewhat surprising given the vast range of linguistic and psychological literature on the subject. Perhaps it is because linguists and psychologists can try to digest just one piece of the acquisition puzzle, while a computational model must typically try to gobble a major chunk of language acquisition whole, or risk being called a mere toy. In this light, Satake should be congratulated for trying to present, in one brief volume, a computational model that attempts to handle facts about morpheme acquisition and intonation; varying word order across languages; verb subcategorization; and classic rule overgeneralization, while at the same time at least paying some attention to what psychologists know about child language.

One then obviously runs the risk of stretching too thin, and in fact the volume under review runs far too short in large type. Readers looking for answers to these rich subjects mentioned just above will come away disappointed by a sketch that ultimately can only approximate what computer modeling did in this area more than ten years ago (work by Anderson 1977; Selfridge 1981; and Berwick 1979, 1985). The book exercises the model with a very limited range of sample sentences—just nine examples, with no recursion. More unfortunately, given the emphasis on free word order, no Japanese examples are included. The first quarter of the book is devoted to a rather thin outline of some of the basic psychological results on input available to the child and learnability theory, while the remainder is devoted to the three components of (sub)category generalization, a case analysis of the system working on the examples and a short study of over-regularization, and the use of teacher correction in a so-called "production mode" to repair mistakes.

This last point is quite important, for Satake's intended novel contribution to this older literature is stated clearly at the outset: to build an *empiricist* model of acquisition that is cognitively faithful—that is, one where the structure of language is "out there" in the world and formed by inductive generalization, special properties of parental input (motherese), the order of examples (including negative examples), and the like rather than "in there"—the child's head. Satake means this of course as the polar opposite of "innate" acquisition procedures, which assume a richly structured knowledge of language to begin with. (Satake labels these as "passive" acquisition models

because in these models presumably the child is somehow not actively constructing hypotheses. However, this distinction, like many in the book, is not entirely accurate, as even those models with considerable given structure must actively pick and choose from a range of hypotheses, and the inductive problems are by no means alleviated; see, e.g., Wexler and Culicover (1980)). Satake's new program, BUD ("Bringing Up a Daughter"), substitutes, as all empiricist approaches must, a model with a rich given knowledge of language with a general induction procedure for building grammars. Aside from an admirable focus on some intonational information to mark sentence and possibly other phrase boundaries and an attempt to deal with nonrigid word order languages (such as Japanese), his acquisition procedure is in fact most closely allied with Selfridge's CHILD program (1981), as Satake notes. One of the strengths of Satake's book is to rightly emphasize that we must now go beyond Selfridge's conception, to handle non-English languages. Regrettably there has been until now very little attempt to model the acquisition of a wide range of languages (but even that now is changing as richer cross-linguistic models have been developed; see Clark, Berwick, and Fong (forthcoming) for applications to German).

The key question, though, is whether the empiricist-oriented BUD succeeds and winds up as really all that different from its supposed opposite. This is by no means clear. We can see this by considering the central properties of all language acquisition models as they apply to Satake's BUD: what input is assumed; how that input is represented; what target languages can be acquired; and what learning procedure is available. In Satake's case, the inputs are strings of words/morphemes, simple intonational patterns marking some key phrase (sentence) boundaries, and, crucially, an associated semantic-relational representation (the usual thematic or case frame representation of who did what to whom, p. 51). Negative evidence is permitted, in contrast to most recent work in the field, as Satake attempts to argue—unconvincingly, in my view—that it should be admitted (for one thing, it makes the class of learnable languages too large, as Gold (1967) noted). Importantly, the child is assumed to already have at its command a conceptual base. Of course, this way of stating things *must* in a sense be true for all language acquisition models, since language is a pairing of sounds and meanings. The question is how the child learns the pairing by paying attention to what is relevant and ignoring what is irrelevant.

So what does BUD learn? Interestingly, only finite linear sequences of surface strings are acquired—that is, no recursion is admitted, this being an extension left for future work (Chapter 9). Subcategorization, thematic role order (agent–action–affected object), and inflectional information are acquired on a word category basis by collapsing similar contexts. Satake justly notes the importance of words and word learning. However, this is also an essential feature of most current nonempiricist models (under the heading of the "lexical learning hypothesis"). Finally, the mapping between word sequences and thematic roles is learned, essentially by exhaustively considering all possible orders. For instance, the child can learn that in *the girl ate an apple, the girl* is the agent and *the apple* is the thing eaten (in contrast to the other way around). There is a "production mode" in which an external teacher can correct the errors that BUD makes—something considered exceedingly unlikely as experts in the field have noted and parents worldwide know (again see Gleitman et al. for a review).

On closer inspection, what is learned really amounts to an acquisition of a primitive X-bar syntax, based on some intonational cues and similar trailing suffixes. Satake needs the notion of the *head* of a phrase, and uses it. For instance, the pattern *the ...*, if followed by both *boy* and *girl*, leads to both being placed in the same distributional category—Harris's analysis (1951). But it has been plain for at least 40 years now that this kind of induction cannot work for anything beyond "flat" structures, that is,

nonhierarchical and finite-state systems—a weakness that Satake admits to (p. 161). Nonetheless, this is still phrasal syntax, if of a rather simple sort. The syntax is hidden in the thematic role frames. The system presumes, among other things, that noun phrase boundaries are somehow available: this is implicit by giving the system the case frame structure or thematic roles for a sentence such as *John gave the guy the ice-cream*. Otherwise, the system could not know that the second *the* was not part of the first noun phrase *the guy* rather than the second, again as Gleitman et al. note. What remains a mystery is the same fact that has plagued all such empiricist models from the start: if acquisition were driven by the outside world in the manner suggested, we would expect to find much more variability in the timing and envelopes of acquisition and sensitivity to the input data than we do. Further, as Gleitman et al. observe, if children used Bloomfieldian inductive generalization methods for categories, then indeed we would see a fluidity in the features and categories of language from generation to generation, but we do not. Languages and children are more stable than that. Under exceedingly limited input and cognitive conditions—e.g., deaf children of non–ASL-fluent caretakers; whole classes of children with IQs in the 40–50 range—acquisition proceeds normally. Assuming sophisticated induction procedures here just misses something (to be fair, as it does with the earlier computer models cited above). Still, given this robustness under variation, it comes as no surprise that whenever we carefully examine computational models of language acquisition we find the same stability in the form of encoded predispositions, somehow. A full-fledged computer-based empiricist challenge to the miracle of language acquisition has yet to appear, just as the miracle of language acquisition for modelers of all stripes still stands.

References

Berwick, Robert (1985). *The Acquisition of Syntactic Knowledge*. Cambridge, MA: The MIT Press.

Berwick, Robert (1979). "Computational analogs of constraints on grammars." *Proceedings of the 18th Annual Meeting of the Association for Computational Linguistics*, Philadelphia, PA, 49–54.

Bloomfield, Leonard (1933). *Language*. New York: Holt, Rinehart, and Winston.

Clark, Robin; Berwick, Robert C.; and Fong, Sandiway (forthcoming). "Acquisition of language by genetic algorithms."

Gleitman, Lila; Gleitman, Henry; Landau, Barbara; and Wanner, Eric (1988). "Where learning begins: Initial representations for language learning." In Newmeyer, Frederick J. (ed.) *Linguistics: The Cambridge Survey, volume II: Language: Psychological and Biological Aspects*. Cambridge, England: Cambridge University Press, 150–193.

Gold, Eric (1967). "Language identification in the limit." *Information and Control*, **10**, 447–474.

Harris, Zellig (1951). *Methods in Structural Linguistics*. Chicago, IL: The University of Chicago Press.

Selfridge, Mallory (1981). "A computer model of child language learning." *Proceedings, 7th International Joint Conference on Artificial Intelligence*, Vancouver, 92–96.

Wexler, Ken, and Culicover, Peter (1980). *A Formal Theory for the Acquisition of Transformational Grammar*. Cambridge, MA: The MIT Press.

Professor *Robert C. Berwick* heads the natural language processing group at the MIT AI Lab, is a long-time member of ACL, and is associate editor of *Computational Linguistics*. Among other things, in the midst of doing research in language acquisition and complexity theory, he has been doing diary studies of his own children's language acquisition, convincing him that we have a long way to go in understanding just about anything. Berwick's address is: MIT AI Lab, 545 Technology Square, Cambridge, MA 02139; e-mail: berwick@ai.mit.edu