PRCLD 27 (1988)
Panel

WHAT IS FUNCTIONALISM?

Elizabeth Bates*
UC San Diego

Brian MacWhinney
Carnegie-Mellon

For the last fifteen years, we have been involved in collaborative research on language acquisition in children and language processing in adults, across a range of structurally and functionally distinct language types (Bates and MacWhinney, 1979, 1982, 1987, in press; MacWhinney, 1987; MacWhinney and Bates, in press). We have brought those findings together within a framework for the study of linguistic performance called the *Competition Model*, a model that is in turn inspired by a broader approach to the study of language called *functionalism*, defined as the belief that *"the forms of natural languages are created, governed, constrained, acquired and used in the service of communicative functions"* (Bates and MacWhinney, 1982). So defined, functionalism is the natural alternative to theories of language that postulate a severe separate between structure and function, and/or theories that attempt to describe and explain structural facts *sui generis*, without reference to the constraints on form that are imposed by the goals of communication and the capabilities and limitations of human information processing.

Although this definition seems sensible enough as stated, it has become sadly clear to us over the years that the term "functionalism" does not communicate very well on its own. It means different things to different people, and worst of all, there seems to be a Straw Man Functionalism out there in the hustings that causes trouble wherever we go. In this short paper, we would like to compare and contrast the principles of Straw Man Functionalism with an approach that is (we believe) much more reasonable and much more likely to succeed. The Straw Man theory can be summarized with the following six beliefs:

(1) **Grammar is a direct reflection of meaning.** That is, we can explain all universal and particular aspects of grammar by uncovering the meanings they convey.

*Portions of this paper are taken from E. Bates & B. MacWhinney (in press), "Functionalism and the Competition Model". In B. MacWhinney & E. Bates (Eds.), *The cross-linguistic study of sentence processing.* New York: Cambridge University Press.

(2) **Grammar is iconic.** That is, grammatical devices "look like" their meanings.

(3) **Mappings from meaning to grammar are one to one.** For every meaning there is one and only one expressive device, and for every device there is one and only one associated interpretation.

(4) **Mappings from meaning to grammar are deterministic.** If the meaning conditions associated with a given grammatical form are met, the form will always be used.

(5) **Functionalism is anti-nativist**. Grammars are a cultural invention, and biological principles are irrelevant to their description and explanation.

(6) **Functionalism is anti-linguistic.** Functionalist theories of human performance will ultimately replace linguistic theory altogether. We will bury Chomsky and all the other generative grammarians with him!

In fact, we do not believe that any of the above six statements are true, and we have never espoused them ourselves. So let us go through these six Straw Beliefs one at a time, and replace each one with a more viable functionalist account.

(1) Grammars reflect the interaction between cognitive content and cognitive processes. We believe that grammars carry out important communicative work. Like individual lexical items, specific grammatical devices (ordering constraints, bound and free morphemes, suprasegmental cues) are associated with meanings and/or communicative goals. But the association is rarely direct. We think it more useful to think of language as a complex, multivectorial problem space. Many different meanings are competing for expression in a linear (i.e. time-delimited) channel. The limits imposed by human information processing (limits of perception, articulation, learning and memory) may

ultimately prove more important than meaning itself in elucidating why grammars come to look the way they do.

In the Competition Model, we have borrowed the term *cue validity* to refer to the information value of a given lexical or grammatical device for any particular meaning or function. The term comes from Gestalt psychology, where it was broadly used to refer to the information structure of some aspect of the environment for any goal or condition that is of interest to the organism. In an ideal world, an ideal animal would behave in perfect accordance with cue validity. But we do not live in an ideal world, and we are not ideal animals. The relationship between meaning and form in language cannot be perfect, because of all the constraints imposed by our information processing system. Our experiments to date have shown that cue validity strongly determines the order of acquisition of cues by children, and the weights that adult speakers attach to the same cues during sentence interpretation. Cue validity also plays a major role in the sentence comprehension and production profiles displayed by brain-damaged adults suffering from severe forms of aphasia. However, there are still many systematic exceptions to this principle. We have been able to account for most of these exceptions by invoking principles of *cue cost,* i.e. the information processing costs associated with the real-time use any given lexical or grammatical cue. For example, cues that are equally informative can vary in their perceivability (e.g. Hungarian accusative case suffixes that follow a strong vowel, compared with the same suffix following a final consonant). This factor will influence that degree to which adults "trust" this particular cue to meaning, the age at which children come to rely on the cue, and the degree of resistance to impairment associated with this particular cue in sentence processing by brain damaged adults. Similarly, cues can vary in the demands they place on memory: "local" cues that can be used as soon as they are encountered (e.g. a nominative case suffix) seem to have an advantage over "long distance" cues that require storage and comparison across a set of discontinuous elements (e.g. subject-verb agreement), even though the two sets of grammatical devices may both point strongly toward the same meaning (e.g. the actor role in a transitive action). A full account of how grammars come to look the way they do, how and when they are acquired by children, will require an analysis of the complex interplay between meaning (cue validity) and information

processing (cue cost). Grammars represent a compromise among these forces, and for this reason, the communicative function of a given grammatical form may be quite opaque.

(2) Symbolic and indexical relations between form and function. Linguistic forms rarely, if ever, resemble their meanings. There are of course a few examples of words that "sound like" the things they stand for (e.g. Bang!), but these are few and far between. It is even more difficult to think of grammatical devices that bear a literal physical resemblance to their meanings. There is of course the apocryphal claim that natural languages prefer basic words orders in which the subject precedes the verb because human beings "naturally" tend to perceive actors before they perceive their actions. This claim is silly enough that it is not worth pursuing. But if grammars do not "look like" their meanings, then what kind of natural cause-and-effect relationship could be said to hold between form and function?

C. S. Peirce (1932) has provided an analysis of sign-referent relations that may be as useful in the study of grammar as it is in the study of single signs. *Icons* are signs that come to stand for their referents because of a literal physical resemblance (e.g. a stylized picture of a cigarette to indicate a smoking zone). *Indices* are another class of "natural" signs that come to stand for their referents not because of a physical resemblance, but because their participation in the same event (e.g. contiguity rather than similarity). For example, smoke can serve as an index to fire because the two are commonly associated in real life. *Symbols* are signs that bear no natural relation to their referents (neither iconic nor indexical); instead, they carry meaning only because of an arbitrary convention, an agreement that was reached by a particularly community of users. As Langacker (1987) has pointed out, most lexical and grammatical signs bear a symbolic relationship to their meanings. Grammatical devices exist in order to carry out communicative work, but the work they do does not determine their form. However, in the domain of grammar there may well be many cases of indexical causality if we keep in mind that grammars are jointly caused by cognitive content *and* cognitive processing.

To offer just one example, consider the relative clause. This device is typically used to identify referents in discourse (e.g. "The man that sold me the car", as opposed to some other man), a functional motive constitutes in itself only a form of symbolic determinism. However, the functions served by a relative clause can also help to determine its shape. Bindings between a referent and its modifier are easier to make if the two are in close proximity. Hence the function of referent-identification is best served if the relative clause is placed near its governing noun phrase, where other modifiers are located. However, this solution usually poses another problem: the relative clause must interrupt a main clause. Such interruption is costly for two reasons. First, because relative clauses are longer than most modifiers, the main clause has to be held open for a rather long time. Second, because relative clauses resemble main clauses in many respects, there is a potential for confusion (e.g. which verb goes with which noun). In principle, this problem could be solved by placing a warning signal at the beginning of a sentence to indicate that "a relative clause will be placed within the following sentence at some point; you guess which point". Although this is a logical possibility, it should be obvious why it would not work very well. It makes much more sense to place the marker *at the point of interruption*, to keep the listener from chasing down some garden path and to help him/her construct and attach the clause right where it belongs (i.e. near the element that it modifies). Finally, insofar as an interruption is already placing quite a burden on the processor, the interruption-marking device had best be kept short and sweet. Hence the functions of the relative clause have an effect not only on the existence of certain devices (symbolic determinism), but also on their position and overall shape (indexical determinism). In neither case is it reasonable to say that the resulting grammatical device "looks like" its meaning!

(3) Mappings between form and function are many-to-many. Grammars can be viewed as a class of solutions to the problem of mapping non-linear meanings onto a highly-constrained linear medium. The universal and culture-specific contents of cognition interact with universal constraints on information processing, creating a complex multivectorial problem space with a finite number of solutions. Natural languages exhaust the set of possible solutions to this mapping problem, and because these solutions represent many competing forces, they

invariably involve many-to-many mappings between form and function (c.f. Karmiloff-Smith, 1979), with correlated meanings riding piggy-back on correlated bits of grammar. No single meaning (however abstract) can be allowed a grammatical monopoly.

The many-to-many nature of grammatical mapping is both a cause and a result of the instability inherent in linguistic systems. In fact, there may be no stable, perfect pathway through the linguistic problem space. As Slobin (1982) has pointed out, many processing constraints stand in direct competition; hence stability in one area may create instability in another. From the listener's point of view, a given linguistic. marker will signal its meaning most efficiently if it is consistent, salient and unique. But from the speaker's point of view, the same linguistic device has to be easy to retrieve and produce. Hence the clear and perceivable markers that evolve for comprehension are often subject to erosion in the service of rapid and efficient speech output. Faced with these competing demands, languages have been known to cycle back and forth across the course of history, from one set of solutions to another Hence we must view grammars as a set of *partial solutions* to the mapping problem, each representing one pathway through the constraints imposed by cognitive content and cognitive processing. No solution is perfect, and each one is constantly subject to change; but every grammar used by a community of human adults and acquired by their children has to meet certain some implicit but implacable limits of tolerance.

(4) Grammatical mappings are inherently probabilistic. Languages differ qualitatively, in the presence or absence of certain linguistic devices (e.g. word order constraints, case-marking), but they also differ quantitatively, in the extent to which the "same" linguistic device is used at all and in the range of functional roles that the "same" linguistic device has come to serve.

We have given a number of examples of quantitative differences between languages throughout our work (see especially papers in MacWhinney and Bates, in press). One particularly important example has to do with the relative strength of word order versus subject-verb agreement as cues to sentence meaning. In English, word order is rigidly preserved; in almost all structures (we will consider a few exceptions

later), the order that is preserved is Subject-Verb-Object or SVO. In Italian, word order can be varied extensively for pragmatic purposes -- a fact that comes as something of a surprise to those who believe that such pragmatic word order variation occurs only in case-inflected languages (i.e. languages with markers on the noun to indicate "who did what to whom"). The following list (from Bates and MacWhinney, in press) illustrates some possible variations in the order of major constituents in Italian, in a hypothetical restaurant conversation. This short conversation (a fake, but quite plausible according to our Italian informants) contains all possible orders of Subject, Verb and Object.

1. **SVO:** **Io mangerei un primo.** (I would eat a first course).

2. **OSV:** **La pastasciutta Franco la prende sempre qui.** (Pasta Franco it orders always here).

3. **VSO:** **Allora, mangio anche io la pastasciutta.** (Well then, am eating also I pasta).

4. **VOS:** **Ha consigliato la lasagna qui Franco, no?** (Has recommended the lasagna here Franco, no?).

5. **OVS:** **No, la lasagna l'ha consigliata Elizabeth.** (No, the lasagna it has recommended Elizabeth).

6. **SOV:** **Allora, io gli spaghetti prendo.** (In that case, I the spaghetti am having).

Some of these require particular intonation patterns to sound exactly right, and some are definitely better with particular grammatical markers like the object clitic. But all these orders can be found in a large enough sample of free speech, and all of them occur at some point in the input received by Italian children (Bates, 1976).

At one level, this discourse serves merely to illustrate a well-known qualitative difference between languages: Italian has word order options that do not exist in English at all. However, this qualitative variation also has quantitative implications. We have now demonstrated in

several different experiments that Italian listeners "trust" word order -- even good old-fashioned Subject-Verb-Object order -- less than their English counterparts. Given a sentence like "The pencil hits the cow", English listeners from ages 2 to 80 have a strong tendency to pick the pencil as the agent/subject. Given the Italian equivalent ("La matita colpisce la vacca"), Italians are much more likely to choose the cow as the agent/subject. Hence a qualitative difference in the availability of word order types has a quantitative effect even on that subset of grammatical structures that both languages share (e.g. SVO order).

Most of our joint research to date has concentrated on sentence comprehension. But we have also uncovered some interesting quantitative differences in the domain of sentence production. For example, Bates and Devescovi (in press) have described some robust differences between Italian and English in the use of relative clauses. The structural options available in the two languages are the same, at least for the set of structures studied by these investigators. In both languages, it is perfectly grammatical to describe a picture of a monkey eating a banana by saying either "A monkey is eating a banana" or "There is a monkey that is eating a banana". However, English speakers typically use the first option; by contrast, Italian speakers describing exactly the same pictures, under the same conditions, are three to five times more likely to produce a relative clause. This cross-linguistic difference in relative clause use is already well-established in children by the age of three, and it tends to persist even in elderly patients who have suffered left-hemisphere damage. How can we capture a quantitative difference between two structures that are equally grammatical from a traditional grammatical perspective? To be sure, there are some differences between the two languages in the range of functions that control these particular forms. In particular, Italians appear to use the relative clause as a kind of topic marker. But in addition to (and perhaps because of) these differences in function, there are also clear processing differences between English and Italian in the "accessibility" of the relative clause. We have uncovered similar statistical differences between Italian and English children in rates of article omission (greater in English children well before the age of 3), and in rates of subject omission (with much higher rates of subject omission in Italian children even in the stage of first word combinations -- Bates, 1976). Some of these differences (e.g. subject omission) are

treated in current linguistic theory in terms of a discrete set of rules or parameters; others (e.g. article omission) receive no treatment in current linguistic theory at all. We think that these early differences in performance can only be captured by assuming that very small children are sensitive to statistical as well as structural facts about the language they are trying to acquire. Function and frequency co-determine the selection of grammatical forms in sentence production, in language use by adults and in language acquisition by children.

Physicists have made their peace with the counter-intuitive predictions of quantum mechanics, and they now accept the premise that the position of a subatomic particle may be unknowable *in the absolute*. Uncertainty lies at the core of the universe; it is not just a byproduct of our imperfect measures. We argue that the human language processor is also probabilistic at its core. In the Competition Model, the adult speaker's knowledge of his native language is represented in a probabilistic form, and probabilities play a fundamental role in the process of language acquisition. The difference between obligatory rules and statistical tendencies is simply a matter of degree. This does not mean that we ignore the powerful laws that separate one language from another. After all, the values "0" and "1" do exist even in a probabilistic system, and an adult native speaker may thus come to know with some certainty that a particular structure is impossible in his or her language. The difference between our characterization of adult knowledge (i.e. "competence to perform") and the characterizations offered in most competence models lies in our ability to capture the many values that fall between 0 and 1. We describe linguistic representations in terms of a complex set of *weighted form-function mappings*, a dynamic knowledge base that is constantly subject to change.

. In a sense, language acquisition can thus be viewed as a process of *meaning driven distributional analysis,* similar in spirit to the approach outlined some time ago by Maratsos (1982). However, the Competition Model also furnishes some non-linear principles that permit us to capture sudden phase transitions, U-shaped functions, and the effect of rare events -- all the phenomena that forced psychologists to abandon the simple linear associative models of American Behaviorism. Many of these discoveries within our model have fallen out of two approaches to the

quantification and formalization of language learning: (a) mathematical modelling of the effects of cues on choice behavior in sentence comprehension (McDonald, 1986; McDonald and MacWhinney, in press), and (b) computer simulations of the learning process (Taraban, McDonald and MacWhinney, in press). For example, we have discovered that cue validity can be operationalized in two ways: *overall cue validity* (the proportion of all the cases in which an interpretation must be made in which a given cue is available and leads to a correct interpretation), and *conflict validity* (the proportion of cases in which one cue competes with another in which the cue in question "wins"). Both these metrics can be calculated objectively from texts of real speech, and used to predict the choice behavior of children and adults in sentence comprehension experiments. Interestingly, we have discovered that overall cue validity drives the early stages of language acquisition; conflict validity (affected primarily by rare cases, particularly those that are encountered in complex discourse) drives the late stages of learning in older children and adults. With these two statistical principles, we can capture abrupt changes in sentence processing strategies that occur as late as 7 - 10 years of age.

Although the Competition Model has been developed on independent grounds (to deal with facts of acquisition and processing across different natural languages), the model in its current form has a great deal in common with a recent movement that is alternatively referred to as *connectionism, neural modelling* and/or *parallel distributed processing* (e.g. Rumelhart, McClelland and the PDP Research Group, 1986; Elman, 1988). It remains to be seen how strong that relationship will be, but we are at least convinced that the tools we share will prove to be exceptionally important in the next era of language acquisition research. Cognitive psychology has proceeded for more than thirty years without an adequate model of learning. Unfortunately, research in language acquisition has done the same. The new focus on learning in "brain-like systems" is a healthy one, whatever its limits may prove to be. And the new tools (i.e. mathematical modelling, multivariate statistics, computer simulation) are bound to lead to progress. Natural languages are so complex that "eyeball analysis" alone can only take us so far -- probably no farther than we have come to date.

(5) Functionalism is biologically plausible. The innateness issue is one of the major sources of anger and misunderstanding in the field of psycholinguistics. We think that much of this misunderstanding comes from a failure to distinguish between *innateness* and *domain-specificity*. The innateness issue has to do with the extent to which human language is determined by the unique biological heritage of our species. But this biological heritage may include many capacities that are not unique to language itself: our large and facile brain, our particular social organization, our protracted infancy, and a variety of unknown factors that may contribute in indirect but very important ways to the problem of mapping universal meanings onto a limited channel, and to the particular solutions that we have found to that problem. Hence the human capacity for language could be both innate and species-specific, and yet involve no mechanisms that evolved specifically and uniquely for language itself. Language could be a new machine constructed entirely out of old parts (Bates, 1979). The universal properties of grammar may be *indirectly innate*, based on interactions among innate categories and processes that are not specific to language. In other words, we believe in the innateness of language, but we are skeptical about the degree of domain-specificity that is required to account for the structure and acquisition of natural languages.

(6) Functionalist claims are made at different levels. Functionalist theories of performance are not in direct competition with any linguistic theory. Different kinds of functionalist claims require different kinds of evidence. This is a point that we have tried to make in several places (notably Bates and MacWhinney, 1982; Bates and MacWhinney, 1987 and in press), but it is sufficiently important that we think it deserves reiterating here. We distinguish four different levels of functionalist claims, ordered from weakest to strongest (in the sense that claims at the higher levels presuppose that claims at the lower levels are true).

Level 1 focusses on the role of cognitive and communicative functions in the evolution of language proper, and the history of individual languages. Claims at Level 1 constitute a kind of linguistic Darwinism, i.e. arguments that functional constraints have played a role in determining the forms that grammars take today. Where did the tiger get his stripes? Why do

grammars have relative clause markers? A great deal of work in functionalist linguistics is of this historical sort, in particular studies of "grammaticization" (e.g. Givon, 1979; Bybee, 1985). Although this work is extremely interesting in its own right, claims at the historical level have no necessary implications for current language use by adults, language acquisition by children, or the proper characterization of grammatical knowledge. Like the large-scale forces that operate to create mountains and rivers across geological time, the forces that operate across many individuals to bring about historical language change may not be detectable (or even operative) in every individual case.

Level 2 is a synchronic variant of Level 1, focussing on the causal relationship between form and function in real-time language use by adult speakers of the language. Much of our own work with adults is of this sort: we manipulate competing and converging sets of grammatical forms as "causes" to see what interpretations our subjects derive; conversely, we manipulate competing and converging meanings in picture and film description, to see what expressive devices our subjects produce to meet these demands. However, even if we could show a perfect cause-and-effect relation in adults, we could not immediately conclude that children are able to perceive or exploit these relations.

Level 3 presupposes but goes beyond Level 2, focussing on the causal role of cognitive and communicative functions in language acquisition by children. The cause-and-effect work of Level 2 must be repeated at every stage of language acquisition, to determine empirically if and when children are sensitive to the form-function correlations available in the adult model. Furthermore (as we noted earlier), we need a well-articulated theory of the learning process, one that can adequately describe, predict and explain the stages that children go through on their way to adult performance.

Finally, Level 4 is reserved for the claim that facts from Levels 1 - 3 play a direct role in the characterization of adult linguistic competence. A variety of competence models of this sort have been proposed within the functionalist tradition, ranging from *Eastern European functionalism* (i.e. the so-called Prague School -- Dezso, 1972; Driven and Fried, 1987; Firbas,

1964; Firth, 1951), *British functionalism* (e.g. Halliday, 1966), the American school of *generative semantics* (e.g. Fillmore, 1968; Chafe, 1971), to more recent proposals that include *cognitive grammar* (Langacker, 1987; Lakoff, 1987), *construction grammar* (Fillmore, 1987), *role and reference grammar* (Foley and Van Valin, 1984), and several other approaches that either retain the simple term "functionalism" or elect to avoid labels altogether (e.g. Dik, 1980; Kuno, 1986; Givon, 1979. For the sake of simplicity, we will refer to these otherwise rather disparate linguistic theories with the single term *functional grammar.* Although functional grammars are not designed to account for real time processing, they are most compatible with highly interactive models of performance, i.e. with models like ours. For obvious reasons, "modular" theories of performance are instead more compatible with "modular" theories of competence, that is, with linguistic theories that emphasize the autonomy of various components and subcomponents of the grammar (c.f. Berwick and Weinberg, 1984; Bresnan, 1982; Pinker, 1984). It is quite possible that there will ultimately be a convergence between some Level 4 version of functional grammar, and the performance model that we have developed to account for data at Levels 1 to 3. But it is also possible, at least in principle, that there may be a rapprochement between a functionalist model of performance and the various rules and representations that have been proposed within the many-times-revised-and-extended school of generative grammar.

In short, we are not anti-linguistic, nor is our work *directly* relevant to any particular class of competence models. We are consumers of linguistic theory, and we have our own bets about which linguistic theory or class of theories will ultimately prevail. But we are much too preoccupied with problems of a different sort to enter into the linguistic fray. This is an exciting new era in language acquisition research, and time is too precious to be wasted on battles that are best waged elsewhere.

References

Bates, E. (1976). *Language and context.* New York: Academic Press.

Bates, E. (1979). *The emergence of symbols.* New York: Academic Press.

Bates, E. & Devescovi, A. (in press). A functionalist approach to sentence production. In B. MacWhinney & E. Bates (Eds.), *The cross-linguistic study of sentence processing.* New York: Cambridge University Press.

Bates, E. & MacWhinney, B. (1979). A functionalist approach to the acquisition of grammar. In E. Ochs and B. Schieffelin (Eds.), *Developmental pragmatics.* New York: Academic Press.

Bates, E. & MacWhinney, B. (1982). A functionalist approach to grammar. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: the state of the art.* New York: Cambridge University Press.

Bates, E. & MacWhinney, B. (1987). Competition, variation and language learning. In B. MacWhinney (Ed.), *Mechanisms of language acquisition.* Hillsdale, N.J.: Lawrence Erlbaum.

Bates, E. & MacWhinney, B. (in press). Functionalism and the Competition Model. In B. MacWhinney & E. Bates (Eds.), *The cross-linguistic study of sentence processing.* New York: Cambridge University Press.

Berwick, R. & Weinberg, A. (1984). *The grammatical basis of linguistic performance.* Cambridge, Mass.: MIT Press.

Bresnan, J., (Ed.). (1982). *The mental representation of grammatical relations.* Cambridge, Mass.: MIT Press.

Bybee, J. (1985). *Morphology: a study of the relation between meaning and form.* Amsterdam: John Benjamins.

Chafe, W. (1971). *Meaning and the structure of language.* Chicago: University of Chicago Press.

Dezso, L. (1972). *Bevetes e mondattani tipologiaba.* (Introduction to linguistic typology). Budapest: TIT Kozpontja.

Dik, S. (1980). *Studies in functional grammar.* New York: Academic Press.

Driven, R. & Fried, V. (1987). *Functionalism in linguistics.* Amsterdam: J. Benjamins, 1987.

Elman, J. (1988). Finding structure in time. *CRL Technical Reports.* La Jolla: Center for Research on Language, University of California, San Diego.

Fillmore, C. (1968). The case for case. In E. Bach & R. Harms (Eds.), *Universals in linguistic theory.* New York: Holt, Rinehart & Winston.

Firbas, J. (1964). On defining the theme in functional sentence. *Travaux linguistique de Prague, 1:* 267 - 280.

Firth, J. (1951). *Modes of meaning.* London: The English Association.

Foley, W. & Van Valin, R. (1984) *Functional syntax and universal grammar.* New York: Cambridge University Press.

Givon, T. (1979) *On understanding grammar.* New York: Academic Press.

Halliday, M. (1966). Notes on transitivity and theme in English: Part 1. *Journal of Linguistics. 2:* 37 - 71.

Karmiloff-Smith, A. (1979). *A functional approach to child language.* New York: Cambridge University Press.

Kuno, S. (1986). *Functional syntax.* Chicago: University of Chicago Press.

Lakoff, G. (1987). *Women, fire and dangerous things.* Chicago: University of Chicago Press.

Langacker, R. (1987). *Foundations of cognitive grammar.* Stanford, Ca.: Stanford University Press.

MacWhinney, B. (1987). The competition model. In B. MacWhinney (Ed.), *Mechanisms of language acquisition.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

MacWhinney, B. & Bates, E. (Eds.), *The cross-linguistic study of sentence processing.* New York: Cambridge University Press.

Maratsos, M. (1982). The child's construction of linguistic categories. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: the state of the art.* New York: Cambridge University Press.

McDonald, J. (1986). The development of sentence comprehension strategies in English and Dutch. *Journal of Experimental Child Psychology, 41:* 317 - 335.

McDonald, J. & MacWhinney, B. (in press). Maximum likelihood models for sentence processing. In B. MacWhinney & E. Bates (Eds.), *The cross-linguistic study of sentence processing.* New York: Cambridge University Press.

Peirce, C. S. (1932). *Language and context: Collected papers.* Cambridge, Mass.: Harvard University Press.

Pinker, S. (1984). *Language learnability and language development.* Cambridge, Mass.: Harvard University Press.

Rumelhart, D., McClelland, J. & the PDP Research Group. (1986). *Parallel distributed processing: explorations in the microstructure of cognition.* Cambridge, Mass.: MIT/Bradford Books.

Slobin, D. I. (1982). Universal and particular in the acquisition of language. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: the state of the art.* New York: Cambridge University Press.

Taraban, R., McDonald, J. & MacWhinney, B. (in press). Category learning in a connectionist model: learning to decline the German definite article. In *Sixteenth annual UWM symposium on linguistic categorization.*