Biological Perspectives on Language

edited by
David Caplan
Andre Roch Lecours
Alan Smith

Errata

Legends to figures in chapter 3

Figure 3.1, p. 42 Baginski, 1871 (from Moutier 1908, p. 34). *D*—principal center for the construction of ideas. *a*—ends of auditory nerve. *b*—center of acoustic perception. *c*—center of elaboration of thought based on sound (Centrum des Klanggedächtnisses). *a′b′c′*—analogous centers for vision. *e*—center of coordination of movement. *f*—motor pathways. Moutier (1908) adds the note that this was the first "schema" ever traced.

Figure 3.2, p. 43 A diagram equivalent to those of Baginski (figure 3.1), Langdon (figure 3.3), and Moeli (figure 3.4). A—auditory specific analysis (for speech). V—visual analysis (for words). I—modality-independent analysis ("ideas"). O—speech output processes.

Figure 3.3, p. 43 Langdon's model (from Moutier 1908, p. 59; Moutier took it from Hughes 1903).

Figure 3.4, p. 44 Moeli, 1891 (from Moutier 1908, p. 53). *B*—ideational center. *spr* (Sprache)—elaboration of speech. *acoust*, *tact*, *Opt*—sensory pathways. *Geh*, *gef*, *ges*—cerebral centers for hearing, touch, and sight. *x*—localization of a lesion supposed to affect the visual aspects of language. *y*—a lesion giving global aphasia ("l'aphasie totale").

Figure 3.5, p. 45 Ballet, 1886 (from Moutier 1908, p. 47). I—center of the intellect. A—auditory center. V—visual center. P—speech. E—writing.

Figure 3.6, p. 45 Grasset, 1896 (from Moutier 1908, p. 51). O—ideational center. A—auditory center. V—visual center. M—verbal motor center. E—verbal graphic center.

Figure 3.18, p. 56 Viable alternatives from current data as to the connectedness of various processes. (left) The current logogen model (cf. figures 3.7 and 3.11). (right) According to Allport and Funnell (1981). V, A, C, and S refer to visual, auditory, conceptual, and speech processes.

Figure 3.19, p. 58 An expanded version of the logogen model, from Morton and Patterson 1980a. On this diagram are indicated the points of breakdown required to account for the performance of a particular deep dyslexic patient.

Figure 3.20, p. 59 A variant formulation of the processes involved in reading, from Shallice 1981. This version is used to analyze in a single diagram nine types of acquired dyslexia.

Legends to figures in chapter 15

Figure 15.1, p. 365 Schematic description of three hypotheses for selectivity of single neurons in the central nervous system. At the sensory input side, a complex stimulus may be, and in fact usually is, coded by many sensory receptors. Similarly, at the motor output, a complex stimulus may potentially affect the discharge of any or all motor neurons. In between, extreme single-unit selectivity (grandmother cells) assumes there is a level where, of many possible cells, only one responds to the complex stimulus. This is the point of maximum selectivity. A completely distributed model (e.g., a Fourier-transform hologram) may have many or all cells change their discharge in response to a sufficiently complex stimulus. The evidence suggests that an intermediate position is correct. Cells show, even at the point of maximum selectivity, modest distribution coupled with considerable single-unit selectivity. A guess would be that perhaps a few percent of cells in a relevant region of cortex might change their discharge when a complex stimulus appeared. From Anderson and Mozer 1981.

Figure 15.2, p. 368 Consider the properties of two sets of N neurons, α and β. Every neuron in α projects to every neuron in β. This drawing has $N=6$ and understates the size and connectivity of the nervous system by several orders of magnitude. From Anderson et al. 1977.

Figure 15.3, p. 371 A prototype dot pattern (P) followed by five examples at various degrees of distortion. Dots were generated on a 512-by-512 array and presented to subjects on a CRT screen. The number refers to the average number of locations moved on the array. A distance of 100 array locations is indicated. See Knapp 1979; Knapp and Anderson 1983.

Figure 15.4, p. 372 Spatial activity pattern assumed to arise from a single dot of a dot pattern. This is an exponential falloff from the central location. The space constant length is indicated. See Knapp and Anderson 1983.

Figure 15.5, p. 373 Sums of four activity patterns corresponding to four displaced dots, such as would arise from four noisy examples of a prototype stored in a distributed memory. For illustration, the four dots were equally spaced from each other and from the prototype location in the center. Notice the enhancement at the prototype location. The space constant length is indicated. The maximum value of the function was always given the same height, so the relative shapes of the curves can be compared.

Figure 15.6, p. 375 A group of neurons feeds back on itself by way of modifiable synapses. Note the possibility of feedback of a cell onto itself. From Anderson et al. 1977.

Figure 15.7, p. 379 A simple example of a two-dimensional brain-state-in-a-box model. The x and y axes correspond to activities in a two-neuron system. Feedback is applied through the feedback matrix, which has eigenvectors pointing toward corners and with eigenvalues as shown. The curved lines passing through the origin are the boundaries of equivalence regions corresponding to one or another corner. Dots are placed on trajectories every five iterations, and the total number of steps required to reach a corner is placed next to the starting point. From J. A. Anderson and J. W. Silverstein, "Reply to Grossberg," *Psychological Review* 85 (1978): 597–603.

Figure 15.8, p. 380 Four sample letter codings for the simulations. These pictures are convenient representations of the 117-element vectors used to code the 26 letters. The left part of the coding for each letter contains 81 elements and corresponds to a modified point-for-point simple mapping. Letters were drawn on a 7-by-7 grid. Positive values are represented by white in the figure, negatives by black, and zero by cross-hatching. Each grid position containing part of the letter was given a positive value. An inhibitory surround was added by placing minus values in nondiagonal adjacent positions. To accommodate the surrounds, the grid was expanded to 9 by 9, the numerical values involved giving rise to the first 81 positions of the vector. The right-hand representations for each letter show the remaining 36 vector elements which were generated by a simple scheme for line detection. Nine 3-by-3 grids were overlapped on the 7-by-7 grid. Each grid was analyzed for line segments at four orientations (0°, 45°, 90°, and 135°), which were located on the small grid. The double-width vertical bar of the J occurred because of the overlap of the small grids. This line was detected by two small grids. From Anderson and Mozer 1981.

Figure 15.9, p. 382 Representations of the five eigenvectors with largest positive eigenvalues plus one additional eigenvector. For display purposes, only significant positive and negative values (> 0.1) are represented as white (positive) and black (negative). In reality, these vectors are continuous-valued. The same coding scheme demonstrated in figure 15.7 is used to draw these representations. These vectors correspond to "macrofeatures." Note the general lack of immediately obvious interpretation. From Anderson and Mozer 1981.

Figure 15.10, p. 385 Results of the numerical ablation study. The initial matrix was 90 percent connected, then matrix elements were removed about 4 percent at a time. The average correlations between the final states of the system with and without the ablation are plotted. A value of 1.00 would mean there was no effect of the ablation in terms of the categorization behavior of the system.

Figure 15.11, p. 386 The number of iterations required to reach a corner in the ablation study. Only 100 iterations were allowed for each initial state. If the state vector was not completely saturated after 100 iterations, that particular vector was not used to compute the average. No corners were reached for any of the 26 letters if no value is given in the graph.

Figure 15.12, p. 392 Vectors representing "words" and "assertions" in the distributed inference model. A plus represents +1 in the stimulus coding, a minus −1. Blanks correspond to 0. The matrix learned these 50-dimensional vectors. The "words" (16-dimensional Walsh functions, chosen for their convenience) were combined into a set of five "assertions," which were learned by the learning matrix.

Figure 15.13, p. 394 For testing the inference abilities of the matrix, a set of "QUERY" stimuli were developed (first set). A query had part of the assertion missing, and the model was required to reconstruct the correct values for the missing part of the vector. In the QUERY stimuli, + represents +1, − represents −1, and a blank is a 0. In the second set, we see the attempt of the matrix to answer the QUERY. These vectors represent final states of the system (i.e., corners) after 100 iterations through the matrix. The limits of saturation in the functioning system were +1.5 and −1.5; thus, + represents +1.5 and − represents −1.5. The QUERY stimulus set is arranged in pairs, with the complete assertion next to a query so the final states can be compared. In 11 out of the 15 cases the final states are identical, in two others they are close (two components different), and in the other two they are very different. The matrix was first taught the words, and then the assertions. All the diagonal matrix elements were set to zero, as were 30 percent of the matrix elements. A small amount of Gaussian noise was added to each presentation of the stimuli during learning. There were 1,500 presentations of the words to form the initial matrix, and then 400 presentations of the assertions.

Figure 15.14, p. 395 A particularly difficult set of stimuli for correlational models to learn. The words and the query set are given, as well as the final classifications when the matrix was formed. Details of the simulation were very similar to the set of assertions involving relations: 30 percent of the matrix elements are zero, the main diagonal is zero, and there was a modest amount of Gaussian noise present during learning. There were 1,500 presentations of the words at first, and then 400 presentations of the assertions.

Series Foreword

The MIT Press Studies in Neuropsychology and Neurolinguistics present theoretical and empirical research on the neural mechanisms underlying language and cognition and their pathologies. These studies focus on the application of concepts and methods of the study of normal language and cognition to the investigation of their disturbances and on the use of increasingly detailed observational techniques in the clinical and basic neurosciences to the neural bases of language and cognition (normal and disturbed). The scope is broad, including humans and nonhumans, development, and related issues. The series is intended to make available significant new work in neuropsychology and neurolinguistics, to disseminate new approaches, ideas, and results, and to catalyze new research.

The members of the advisory board for this series are Maureen Dennis of The Hospital for Sick Children in Toronto, Patricia Goldman-Rakic of Yale University, Steven Hilliard of the University of California at San Diego, Marc Jeannerod of INSERM at Lyon, Scott Kelso of Haskins Laboratories, André Roch Lecours of the University of Montreal, John C. Marshall of Oxford University, Larry Squires of the University of California at San Diego, and Edgar B. Zurif of the City University of New York and Brandeis University.

David Caplan

Chapter 1

The Mental Organ for Language

David Caplan

One of the most productive concepts in contemporary psychology is that cognitive function is the result of a number of autonomous processes, each operating in a particular intellectual sphere, whose individual natures and interactions determine the cognitive capacity of an organism and thereby influence its behavioral repertoire. One area of cognition in which this concept has emerged most clearly is that of human linguistic abilities. Investigators have spoken of the existence of a "language-responsible cognitive structure" (Klein 1977) and a "mental organ for language" (Chomsky 1977)—terms intended to imply that the human capacity to utilize language is one of the functional entities that make up human mental endowment. I shall consider some of the issues raised by the notion of an autonomous capacity for language and illustrate some of the properties of language structure and processing that characterize this hypothesized autonomous component of the human mind.

Among the features of language which investigators have recognized and which have led them to postulate a mental organ for language are the following three:

- Language as we conceive of it pretheoretically is unique to humans. Though many animals have systems of signal and sign display that allow for restricted types of communication, for expression of individual, social, and species identity, and for other functions served by language, our impressions are that what humans call language is not found in comparable form in other species, and that human language and other systems differ in qualitative rather than simply quantitative ways. The qualitative differences are felt to exist at the levels of formal structure, semantic power, and social and personal utility.
- Language is universally present in humans. Except in cases of neurological impairment and failure to be exposed to linguistic environments,

all humans develop language in auditory-oral modes. Humans born deaf and raised in communities and families of the deaf develop language in manual-visual modes, and the language they develop is, within limits imposed by the mode of expression, highly similar in form and probably in semantic power to the auditory-oral language that hearing humans achieve.

- Humans learn language without explicit instruction. In comparison with the highly structured input that characterizes educational efforts in other domains (consider simple arithmetic), linguistic input to children is quite unstructured and poorly reinforced. Furthermore, children appear to acquire language in regular sequences, despite diverse exposure conditions and sequences of exposure, and they acquire knowledge about language that goes well beyond mastery of the actual exemplars to which they are exposed. It would appear, therefore, that language acquisition does not utilize environmental input in the same way as is the case in other domains in which "learning" takes place.

These three observations ("claims" might be a better word) are central premises on which the concept of a mental organ for language has been based. I shall not review the data that support them; rather, I accept them as a point of departure for the characterization of a postulated mental organ for language. We can conceive of an explanation for these three observations if we postulate that humans are naturally equipped to develop and utilize what we call language, that this "being equipped" is an inherent capacity of humans that develops with the maturing nervous system, and that other species lack this "equipment." Klein (1977) presents four postulates that spell out this position more explicitly:

- What accounts for a person's capacity for language is that he possesses a language-responsible cognitive structure. Although the structure of the LRCS may vary somewhat from speaker to speaker, certain universal features are common to the LRCS of each normal individual.
- Considerable uniformity exists in the onset time and the course of language acquisition among normal individuals. These universal features of language acquisition are controlled to a significant extent by maturational changes in the LRCS.
- The physical realization of the LRCS is a brain structure. This realization is sufficiently uniform among neurologically normal individuals that an (idealized) description of the LRCS as a neural structure is possible.
- The possession of the LRCS is a species-specific trait; that is, universal features of the LRCS (both structural and developmental) are con-

trolled by genetic factors characteristic of every normal member of the human species.

These postulates, which embody much current thinking about the biology of language, present a general characterization of the concept of a mental organ for language that links the view of the species-specific cognitive ability to a particular material basis (a neurological "realization" based on a genetic endowment) and to a strongly nativist psychology.

Accepting this characterization as a point of departure, I propose to focus on work that provides a more detailed characterization of the postulated mental organ for language. I shall emphasize the issue of the autonomy of linguistic systems from other cognitive systems. Evidence that language faculties constitute an autonomous aspect of human mental life would be, *ipso facto*, evidence that, if there is a mental organ in the sense discussed by Klein, it has particular highly restricted domain of application—namely, what we might intuitively appreciate and ultimately theoretically define to be human language.

Moreover, the notion of the domain specificity, or autonomy, of a mental organ for language invites us to phrase specific hypotheses about phylogeny and ontogeny. Phylogenetically, we might consider that what separates humans from other species is our possession of just this particular domain-specific mental capacity, at least in its entirety. We might look for separate evolutionary lines, as in the vocal learning seen in birdsong, or intelligent behavior we can discover by observation or experimentation in other primates, and compare the resulting functional abilities with those of language. It may turn out to be the case that, when characterized in detail, various aspects of the behavioral repertoires of other species are in fact highly similar to parts of a language faculty. This might provide us with animal models for the physical basis of at least parts of a language faculty. Ontogenetically, we could consider that what it is that develops under genetic control of neurological maturation, with appropriate environmental exposure, is just the restricted cognitive capacity that we associate with language. Other aspects of the ontogeny of cognition and intelligent behavior could be the result of very different psychological processes and might follow quite different principles of development.

What, then, are some of the features of a postulated mental organ for language? Current linguistic and psycholinguistic work suggests that we conceive of a mental organ for language as consisting of two separate related elements: a set of knowledge structures and a set of procedures that utilize these knowledge structures. I shall deal with each of these features in turn.

There are several features of linguistic knowledge that have been emphasized in the modern literature and that we may take as ontological claims about this aspect of cognitive psychology. Some are psychological, such as the observations that, though we are aware of certain aspects of the structure of language (such as the existence of discrete elements which we term words), much linguistic knowledge is unconscious, or the claim that important aspects of this knowledge are innate and serve as a basis of the child's ability to identify and assimilate features of the particular language to which it is exposed. Some of these features are structural; "core" linguistic knowledge is highly detailed, has structural features not found in other domains of the intellect, and is itself modular in the sense that it consists of structural components, each representationally distinct and coherent, which interact in highly constrained ways to yield the full array of structures that make up linguistic knowledge. These latter structural features—modularity, uniqueness, and representational coherence—are those that lead to the conclusion that, at the level of knowledge structures, a mental organ for language is an autonomous aspect of human intelligence.

To exemplify one analysis that presents hypotheses about the type of knowledge that might be represented in a mental organ for language, we may consider some quite simple and uncontroversial facts about English syntax, the analysis of which has provided some of the most interesting contemporary hypotheses about certain aspects of language structure. The phenomena in question are restrictions on the possible movement of constituents of sentences in English, a domain of empirical observation since the earliest days of transformational theory. The following analysis is highly simplified and omits many discrepant and technical details, but it does illustrate at least some of the basic properties of the knowledge structures in a postulated mental organ for language.

Among the data to which any theory of grammar must be responsive are facts about English such as that only the first of the three following sentences is acceptable.

1. The boat that you believe John painted is a yawl.
2. *The boat that you believe the claim John painted is a yawl.
3. *The boat that you asked who painted is a yawl.

It is clear that sentences 2 and 3 are not ruled out on semantic grounds; if they were well formed, their meanings would be perfectly clear. They must be unacceptable because of their syntactic structure.

Simplifying considerably, we may say that sentences 2 and 3 are unacceptable in English because the noun phrase *the boat*, which serves as the subject of the principal verb of the main clause, is also the direct object of the verb in the embedded clause and cannot serve both these functions simultaneously in structures such as those found in sentences 2 and 3. Given the acceptability of sentence 1, it cannot simply be the case that English does not allow one noun phrase to fulfill these two thematic functions simultaneously; rather, there must be some feature of English grammar that is relevant to the distinction between the first sentence and the other two and that is derived from the syntactic structures of sentences in which a single noun phrase fulfills both these thematic roles.

It minimally follows that in order to be able to describe and ultimately explain facts such as these we need a system for representing the relevant aspects of the structure of English sentences. A wide variety of such descriptive frameworks have been proposed; I shall present one in a very simplified way. This system postulates that English contains structures called clauses, which themselves consist of an introductory complementizer and a propositional content (Bresnan 1972). We may assign each of these elements distinctive labels, using the symbol S (which we may think of as sentence) for the propositional content of a clause, and the symbol $\bar{S}$ for the entire clause, including its complementizer. In other words, we are assuming that the basic rules of English contain the following two phrase-structure rules:

4. $\bar{S} \rightarrow \text{COMP} + \text{S}$
5. $\text{S} \rightarrow \text{NP} + \text{VP}$.

Within this framework, movement rules place a subset of words (*who, which, that*, and so on) in a complementizer position.

We are now in a position to assign highly simplified syntactic structures to the relevant portions of sentences 1–3:

6. The boat [$_{\bar{S}}$[$_{COMP}$ which that][$_S$ you believe [$_{\bar{S}}$[$_{COMP}$ that][$_S$ John painted t]] is a yawl.
7. The boat [$_{\bar{S}}$[$_{COMP}$ which that][$_S$ you believe [$_{NP}$ the claim [$_{\bar{S}}$[$_{COMP}$ that] [$_S$ John painted t]]]]] is a yawl.
8. The boat [$_{\bar{S}}$[$_{COMP}$ which that][$_S$ you asked [$_{\bar{S}}$[$_{COMP}$ who][$_S$ t_1 painted t_2]]]] is a yawl.

The symbol t stands for trace and marks the original position in the embedded clause of the noun phrase *the boat*, which has been moved to the complementizer position and replaced with the word *which*. The resulting complex complementizer *which that* in sentences 6–8 is ultimately reduced to the single complementizer *that* in sentences 1–3.

Given syntactic structures with this level of detail, we are in a position to identify the differences between sentence 1 and sentences 2 and 3. This statement requires that we identify the symbol called a bounding node, which has the property of constraining movement rules such as the one that moves *the boat* to the complementizer position. The elements in movement rules cannot be separated by more than one bounding node (Chomsky 1977). There is evidence that the node NP is the bounding node in all languages. In addition, most languages have a second bounding node that relates either to clauses or to propositions. In English the bounding node is related to the propositional content of the clause, S. In sentence 1 (structure 6), *which* moves to each COMP in turn and is never more than one bounding node from its previous position. In sentence 3 (structure 8), the lower COMP is filled, so *which* must move directly to the higher COMP. In sentence 2 (structure 7), the lower COMP is separated from the higher COMP by both NP and S. Thus, *which* would have to move over two bounding nodes at some point in the generation of sentences 2 and 3. Thus, sentences 2 and 3 violate the conditions on movement transformations in English.

Other languages appear to have different choices of bounding nodes. Rizzi (1978) has argued that in Italian the bounding node is $\bar{S}$ rather than S, a situation that leads to the acceptability of sentences 1 and 3 and the unacceptability of sentence 2. Moreover, the acceptability or unacceptability of structures such as 1–3 is critically dependent on the fact that the constituent noun phrase, *the boat*, is moved without leaving an overt marker in its original position. There are languages (such as Hebrew) in which such a marker is found in the form of a resumptive pronoun, a different process than movement; in such languages, structures such as 2 and 3, with the appropriate resumptive pronoun in the embedded clause, are acceptable.

The analysis of constraints on movement transformations, and many other analyses couched in a similar vocabulary, have suggested properties of a mental organ for language. It seems clear that the types of knowledge structures that are a prerequisite for the statement of conditions on movement transformations are unconscious aspects of our knowledge of English. It is clear that children receive no specific instruction about such structures and constraints, and it seems highly improbable that any general learning strategy would arrive at an appreciation of such structures. On the contrary, it has been argued that the process of language acquisition is understandable only if we postulate that the child comes to the identification and analysis of certain "core" features of sentences such as 1 equipped with a rich conceptual apparatus in which only a small number of parameters, such as the choice of S or $\bar{S}$ as bounding node, need to be specified (Chomsky 1977, 1982).

This analysis, and others like it, may also serve to illustrate and justify the concepts of the autonomy and the modularity of linguistic structures.

The issue of autonomy of the representations that constitute knowledge structures in a postulated mental organ for language involves the uniqueness and internal coherence of the formal representations needed to express linguistic knowledge. It is quite clear that contemporary cognitive psychology does not make use of substantive symbols or of the organization of such symbols, which appear in the description and explanation of linguistic regularities and which, by hypothesis, are the knowledge structures represented in a mental organ for language. The statement of conditions on movement transformations that I have just presented, even in this schematic form, is possible only within a very highly articulated framework of syntactic structure, which does not appear in existing theories of other cognitive abilities. The mental organ for language would thus appear to be separate from other faculties with respect to its particular knowledge structures, in the sense that it uses structures not found elsewhere.

The converse of this statement is less clear. There do seem to be plausible candidates for overlap between structures found in theories of other aspects of mental life and linguistic representation. For instance, a variety of logical notations have found a role in capturing linguistic regularities as well as in describing human reasoning, prototype theory seems to apply to linguistic as well as nonlinguistic cognitive categories, and the devices of recursion and transformation utilized in many versions of syntax are borrowed from mathematical formalisms and may find a place in other theories of cognition. Though a certain degree of caution is indicated in accepting the conclusion that these substantive and formal devices are identical in the domains of linguistic and nonlinguistic knowledge, we must certainly accept the possibility that language may borrow from a variety of formal and substantive sources and achieve representational autonomy only in the sense that it adds to the formalisms utilized elsewhere and not in the sense that none of its formal and substantive devices have parallels in other systems.

The final feature of linguistic representations that I wish to emphasize is their modularity. I have noted above that the statement of conditions on movement transformations requires a certain set of syntactic representations. It is equally important in the present context that it does not require other aspects of linguistic information. Thus, for instance, although some forms of lexical information (that is, information that depends on particular words) can influence phenomena such as the conditions on movement and on other transformations, there seems to be no case in which the actual phonological sound of a word is relevant to the statement of such conditions. In fact, the sound pattern of individual words is a highly complex linguistic structure. The fact that it does not influence other complex phenomena in natural language suggests that each of these phenomena depends on a separate system, and that these systems combine to yield the totality of linguistic structures only by the input and output to and from various subsystems of linguistic representations.

It should be borne in mind that the syntactical example I have chosen is but one of a larger number of similar examples of the autonomy and the modularity of linguistic representations. Though the study of syntax is the area in which these features of linguistic representations have been most frequently emphasized with respect to their implications for a postulated mental organ for language, exactly the same points emerge from studies of thematic relations (Bresnan 1982), metrical structures in phonology (Liberman and Prince 1977), morphology (Leiber 1980), and many other areas. If there is a mental organ for language, and if it does contain knowledge structures of the sorts I have very sketchily outlined here, it is indeed a rich system.

The second part of a postulated mental organ for language consists of a set of procedures that utilize linguistic structures of the sort I have sketched. Here, as in the area of knowledge structures themselves, the question of a mental organ for language is closely related to the question of the autonomy of these procedures. In very general terms, what we would like to know is whether the procedures that utilize linguistic knowledge are specific to this purpose or whether they are particular applications of mental procedures that operate on other sorts of mental representations. It is far beyond the scope of this chapter to attempt a plausible characterization and taxonomy of psychological procedures, and I shall again present simply one example of a psycholinguistic process (sentence parsing) for which the delineation of general features of the process itself and analyses bearing on its particular nature are fairly well developed and directly relevant. Like my presentation of the analysis of syntactic structures relevant to conditions on movement transformations, the example that follows omits much conflicting evidence and technical detail but, I hope, will serve the purpose of illustration.

Sentence parsing, the assignment of a syntactic structural description to an utterance, is increasingly being understood as one of a number of so-called "on-line" tasks that involve linguistic representations. A number of general characteristics of on-line processes have been suggested (Marslen-Wilson and Tyler 1980). On-line processes are unconscious, not permeable by systems of belief (Pylyshyn 1981), rapid, obligatory, and dependent on information in the physical signal ("bottom-up" in this particular sense), and they interact with other processes in constrained and efficient ways. Sentence parsing falls into the class of on-line processes when seen in these terms. The fragment of an analysis of a sentence parser that I shall present is due to Fodor and Frazier (see Frazier and Fodor 1978 and Fodor and Frazier 1980) and illustrates some of the specific properties that a sentence parser may have.

Consider the following sentences:

9. John bought the book for Susan.
10. John bought the book that I had been trying to obtain for Susan.

One's first interpretation of sentence 9 is that Susan is who John bought the book for, and one's first interpretation of sentence 10 is that what John bought was the book that I had been trying to obtain for Susan. On reflection, it is clear that both sentences are ambiguous; sentence 9 can mean that what John bought was the book for Susan and sentence 10 can mean that it was for Susan that John bought the book that I had been trying to obtain. These possibilities of interpretation, however, are clearly not those that suggest themselves in the first instance. This observation of the relative availability of these two interpretations of the sentence suggests that the parsing of these sentences, in the immediate, obligatory, rapid, unconscious manner that characterizes online psychological processes, arrives at the preferred interpretation before the second. This, in turn, suggests certain features of the human parser. Again, to see what these features are, we need a system of representation that will allow us to capture the differences between the two possible intrepretations of these sentences.

Figures 1.1 and 1.2 indicate a variety of differences between the two structures that underline the two possible interpretations of sentences

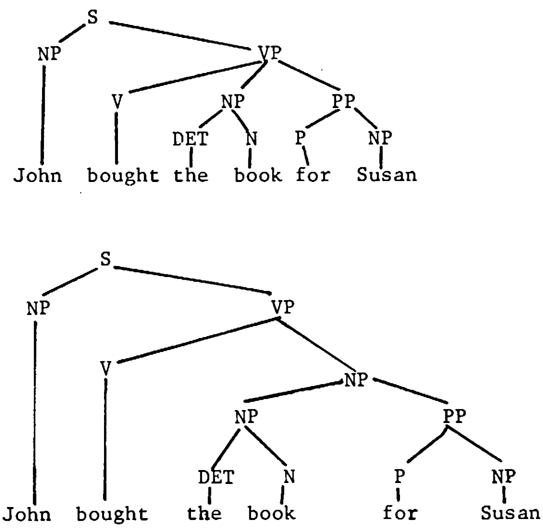

Figure 1.1

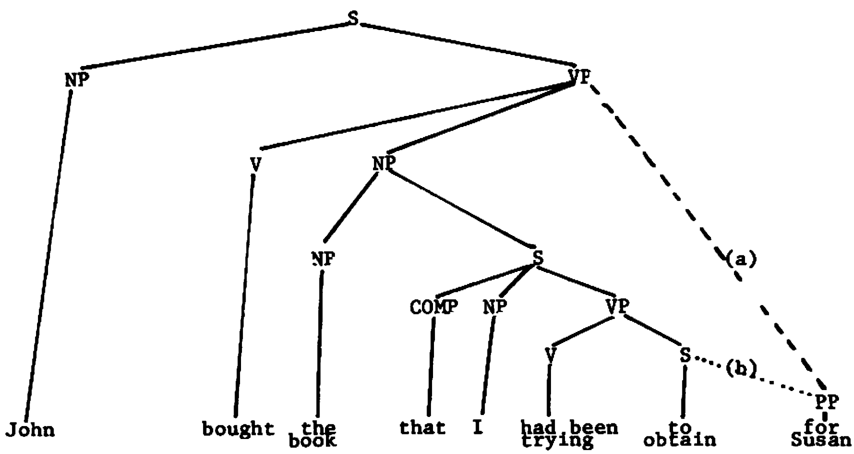

Figure 1.2

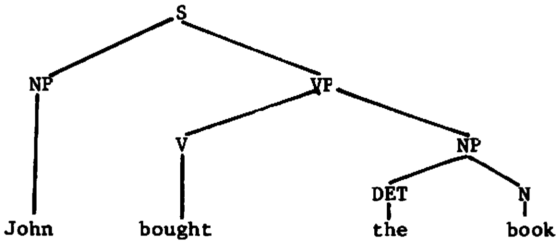

Figure 1.3

9 and 10. Following Frazier and Fodor, we see that the relevant aspect of these structures is the structural configuration into which the node PP is inserted in both these sentences. In the upper diagram in figure 1.1 the prepositional phrase *for Susan* is attached directly to the verb phrase, whereas in the lower diagram it is attached to an NP that is itself attached to the VP. In figure 1.2 the VP *for Susan* is attached directly to a VP in both case a and case b and does not in either case make this attachment by an additional node. This structural difference, in this model, is critically important in deciding which is the preferred interpretation of these two sentences. One aspect of this model postulates that, when a structure (such as an NP) is identified in a sentence, it is attached to the phrase marker that has previously been constructed using the smallest number of nonterminal nodes—a principle called minimal attachment (MA). Minimal attachment comes into play not with the prepositional phrase *for Susan* in figure 1.1, but rather with the previous noun phrase *the book.* Having identified *the book* as an NP, the parser, following minimal attachment, will attach it directly under the VP and will not postulate the second NP node intermediate between the VP dominating *bought* and the NP dominating *the book.* At the point where *Susan* is recognized as a PP, the phrase marker constructed by a parser will have the form seen in figure 1.3. At that point, there is only one place for the PP *for Susan* to be attached: as a "sister" to the NP *the book.* Attaching the PP *for Susan* as in the lower diagram in figure 1.1 would entail both adding an additional node and revising the previously constructed phrase marker—something a very natural principle of parsing would seek to avoid. Therefore, the parser proceeds to attach the PP *for Susan* to the VP, with the result that the sentence is interpreted to mean that it was for Susan that John bought the book.

In the case of sentence 10, the partial structure constructed by the parser at the point where the PP *for Susan* is recognized is illustrated in figure 1.4. As can be seen, attachment of the PP *for Susan* to the VP dominating *bought* and to the VP dominating *obtain* are equally possible according to the principles thus far mentioned; neither attachment requires more nodes than the other and neither requires the revision of a previously constructed structure. What needs to be explained in sentence 10 is why the preferred and immediate interpretation of the sentence is the one in which the PP *for Susan* is attached to the VP dominating *obtain*, rather than there being a natural ambiguity which is immediately appreciated in these sentences. Here the answer lies in a third principle, which Frazier and Fodor term right attachment (RA): Terminal symbols optimally are attached to the lowest nonterminal node.

The phenomena can become considerably more complex. These principles are embedded within a two-stage model of parsing in the Frazier and Fodor proposals that leads to more detailed predictions about preferred and possible interpretations of other sentences. It should be noted that there are other models [in particular, augmented transitional network models (Wanner 1980)] that provide alternate analyses of these and other phenomena related to parsing. Nonetheless, this simplified example will serve to illustrate several important features of the human parser, itself taken as one aspect of the psychological processes utilizing linguistic representations.

As in the case of linguistic representations (and their correlates, the knowledge structures contained in a postulated mental organ for language), this analysis of the process of human parsing raises the issue of autonomy.

The first issue for autonomy is uniqueness. In the case of psychological procedures, we are interested in whether the substantive elements to which procedures apply, and the formal operations and organization of the procedures in this domain of human psychology, contain unique elements. Even this cursory and superficial presentation of some of the features of the human parser strongly suggests that the answer to these questions is affirmative; this analysis suggests that the parser makes use of substantive elements (certain types of linguistic representations) and consists of operations and constraints that, so far at least, have not found parallels in other areas of human psychology.

The second aspect of the autonomy issue with respect to psychological procedures utilizing linguistic information is whether they utilize portions