85

e Occasion
niversary of
nunication by
Clarke,
Dertouzos,
Halle,
ola Pool,
B. Wiesner

n D. deButts
homas E. Bolger

eration with the
Telegraph Company

THE TELEPHONE'S FIRST CENTURY— AND BEYOND

THOMAS Y. CROWELL COMPANY
Established 1834 / New York

SNER

ng plant capacity may
le, it would cost many
to create enough syn-
he present natural pe-
tions may restrict the
ations services, includ-
n expand.
or even the possibility
ajor inhibitions to the
in some areas—for
of health areas, in the
l for the generation of

problem for the next
to be how to assure
In any event, I doubt
future dislocations as
se caused by the shift
and cultural changes
nt of mass-production
hnology including the
radio and television,
medicine and health
Further, the changes
t about by technology
unications and trade
that must ultimately
aw the people in the
de.
t how to direct tech-
ect change so that the
ying. And as we cele-
hone, the question is,
logy and the informa-
rection?

LANGUAGE AND COMMUNICATION

Morris Halle

Morris Halle was born in Latvia, studied engineering at City College in New York and the University of Chicago, then went to Columbia to study with the great linguist Roman Jakobson. In 1949, when Jakobson left Columbia for Harvard, Halle went with him. Halle obtained his Ph.D. from Harvard in Slavic Linguistics and then became Assistant Professor in the Department of Modern Languages at M.I.T. and a member of the staff of the Research Laboratory of Electronics in 1951. He, along with his colleagues, was instrumental in developing M.I.T.'s graduate program in linguistics and will be the first chairman of the new combined department of linguistics and philosophy in 1977. He has won many honors, including election to the presidency of the Linguistics Society. Among his many books are Preliminaries to Speech Analysis *and* The Sound Pattern of Russian.

Questions of language are a most appropriate topic for discussion on the occasion marking the centennial of the telephone, for not only is the transmission of spoken language the primary function of the telephone, but also as I shall have occasion to remark later, the telephone's inventor, Alexander Graham Bell, had a deep scientific interest in problems of language and must be credited with an important conceptual advance in our understanding of the sound structure of language.

Since early in 1975, M.I.T. has been host to a series of workshops on problems of language and cognition, which were organized with the help of a grant from the American Telephone and Telegraph Company as one of several such activities leading up to a two-day convocation. The workshops brought together between 20 and 30 workers from different research groups with significant involvements in problems of language. Our thinking, like that of most students of language in this generation, had been influenced by the work of our M.I.T. colleague, Noam Chomsky. Of particular significance to our enterprise was Chomsky's insistence that the proper and overriding aim of linguistic description must be to provide an account of the knowledge that native speakers have of their language. The importance of this focusing on the cognitive aspect of language may not be immediately apparent to the nonspecialist. As a participant-observer of the history of the field during the last quarter century, however, I have no doubt that few things have had such a far-ranging effect on the development of the entire field as Chom-

sky's insistence that language is a form of knowledge peculiarly accessible to humans and that it is, therefore, akin to other manifestations of our cognitive faculties such as our ability to perform computations, to play games and to invent, plan and execute complicated structures and maneuvers.

If language is knowledge that is peculiarly accessible to humans, then all language must share certain essential features which are especially well matched to the intellectual capabilities of the young child, for in all known linguistic communities command of language is acquired at a very early age. That languages share many substantive properties is by now a well-established fact: every known language forms sentences by concatenations of words; in every language ever studied, words are made of sequences of a restricted number of speech sounds, and these in turn, as Bell was one of the first to see clearly, are complexes of a small number of phonetic properties, and so forth. In addition to such substantive universals, languages also share properties of a more abstract kind, which in the linguistic literature are now designated by the term *formal universals.*

It was one of Chomsky's suggestions that an important formal universal of language is the syntactic transformation, a special computational device first described by Zellig Harris. A large portion of Chomsky's work as well as that of others during the last twenty years has been concerned with establishing the character and the proper role of the syntactic transformation. Initially transformations were assigned a very large role in the functioning of the language. More recently it has become evident that transformations were unsuited to some of the tasks that were assigned to them. This discovery, as Joan Bresnan has noted, elicited different responses from different researchers. Some proposed to overcome the difficulties by increas-

ing the power of transformations; whereas others have followed the lead of Chomsky and opted for a significant limitation on the power and role of transformations. Though participants in the workshop tended to agree with the latter view, there was no consensus about the precise character of the limitations that should be imposed on transformations. A proposal that seems to hold considerable promise has been presented by Bresnan.

Among the tasks that transformations were especially unsuited for was the characterization of relationships among words. Since relationships among words obviously play a big role in a person's knowledge of language, other means will have to be found to express them. The obvious candidate for expressing this information is the lexicon, for it is there that the speaker's knowledge of the words of his language is represented. The information about the relations that a given word bears to any of the others will then be part of that word's entry in the lexicon. The complexity of this information requires a fundamental revision of our conception of the character of the lexicon. Rather than being a simple listing of more or less odd facts, the lexicon must now be regarded as an active device possessing a structure that must be carefully investigated.

By a most fortunate coincidence, renewed interest in the lexicon developed at about the same time also among psychologists. George Miller and Philip Johnson-Laird have just completed their monumental *Language and Perception,* a large portion of which is devoted to an inquiry into the form and function of lexical entries. Miller's description of this work generated considerable excitement as it dealt with topics that had been uppermost in the minds of several of us. It also led to a very fruitful exchange of information about

much unpublished work that had been going on here as well as elsewhere. One side product of this was that Ray Jackendoff started to come to our meetings and to take an active part in our discussions. The main results of those discussions are contained in the papers prepared by Miller and Jackendoff.

The issue that provoked the most heated debate among us was the manner in which knowledge of language is utilized in the production and understanding of utterances. While few would disagree with Chomsky's remark that "a reasonable model of language use will incorporate, as a basic component . . . the speaker-hearer's knowledge of the language," there was wide divergence about the precise way in which a processing model of language incorporates this knowledge. The problem might perhaps be clarified by an example, which I have adapted from George Miller's paper. Miller points out that the words *woman* and *person* both designate individuals that are animate and human; they differ in that the word *woman,* in addition, indicates that the individual is female, whereas the word *person* provides no information about the individual's sex. If, as appears plausible, the semantic information just presented has to be used by the speaker in the understanding of sentences, then it might be supposed that the word *woman* is semantically more complex than the word *person,* and this difference in complexity would be reflected by differences in the time it takes to understand such otherwise identical sentences as *John's wife is the woman on the right* vs. *John's wife is the person on the right.* It appears, however, that no such differences have ever been found.

What is one to conclude from this negative result? Some would conclude that this result suggests that the postulated difference in semantic content of the words

woman and *person* lacks psychological reality. Others have argued that since the difference between the two words can be established by other behavioral tests, the lack of difference in reaction time does not speak to the issue of psychological reality. While with respect to the example in question the evidence appears strongly to favor the latter view over the former, the evidence is less clear with respect to other aspects of linguistic knowledge. There are particular disagreements about the psychological reality of transformations. On the one hand, Eric Wanner and Ronald Kaplan propose that the transformational model should be replaced by a nontransformational model of their own devising. On the other hand, Merrill Garrett and Kenneth Forster argue that the evidence favors a processing model that incorporates a transformational model of the more familiar kind. It is obvious that considerable further work and thinking will be required before a consensus can emerge.

Last but by no means least, the workshops considered also the insights into the nature of language that might be hoped for from observations of language acquisition by young children and of language loss in the brain-injured. Details of these discussions were presented by Susan Carey, by Michael Maratsos, and by Sheila Blumstein, Mary-Louise Kean, and Edgar Zurif. Because of the extreme variety of the data that have been amassed here, great difficulties are encountered in attempting to interpret them properly, and agreement concerning the significance of a given observation is often impossible to achieve. This difficulty, however, does not detract from the importance of these facts and there are among them many observations that speak loud and clear. Consider, for example, the fact pointed out by Susan Carey that many six-year-olds have a vocabulary of 14,000 words or more. As Carey

noted, this means that the child must be learning words at a rate of about one an hour for every waking hour of early childhood. To make this rate of learning at all plausible, it must be assumed that much of the vocabulary is acquired after an extremely small number of exposures, perhaps no more than one or two. Where this leaves the complicated reinforcement schedules and elaborate learning strategies that make up such a large part of the literature on learning is a question that surely bears looking into seriously.

While all these studies speak to the fundamental aim of our search to gain a better understanding of the knowledge that native speakers have of their language, the extremely compressed fashion with which of necessity I had to deal with the different contributions of my colleagues has not permitted me to convey to you as clearly as I should like to the nature of the object that we are studying and the extent to which we have a grasp of it. In my experience the best way of providing this kind of insight is by examining a number of real examples in detail. And that is what I propose to do next. The examples that I shall be discussing are taken from the phonic aspect of language. I have chosen these because, on the one hand, this is the facet of language that I am most familiar with, and, on the other hand, because this allows me to bring in Alexander Graham Bell and his contribution to the study of language.

I want to begin by observing that the native speaker of a language knows a great deal about his language that he was never taught. As an example of this untaught knowledge, a list appears below of a number of words chosen from different languages including English. In order to make this a fair test, the English words in my list are words that are unlikely to be fa-

miliar to the general public, including most crossword puzzle fans.

ptak thole hlad plast snam mgla vlas flitch dnom rtut

If I now were to ask for a show of hands on each of these 10 words as to whether or not it is to be found in the unabridged Webster's, I am reasonably sure that the majority would vote that *thole, plast* and *flitch* are English words, whereas the rest are not English. This evidently gives rise to a question: Since you have never seen any of the words in the list, how do you know that some are English and others are not? The answer is words judged not English have letter sequences not found in English. This implies that in learning the words of English the normal speaker acquires also knowledge about the structure of the words. The curious thing about this knowledge is that it is acquired although it is never taught, for English-speaking parents—I can swear that this is true of my wife and me as well as some of our acquaintances—do not normally draw their children's attention to the fact that consonant sequences that begin English words are subject to certain restrictions which exclude words such as *ptak, snam* and *rtut*, but allow *thole, flitch* and *plast*. Nonetheless, in the absence of any overt teaching, speakers acquire this knowledge somehow, and this is surely a puzzle worthy of the attention of some learning theorist.

In order to get some insight into how humans acquire knowledge about their language without being taught, it is necessary to understand the character of the knowledge that is being acquired. It is obvious that in the example under discussion that knowledge being acquired concerns the sounds and sound sequences found in English.

Linguists have special ways of dealing with sounds which, incidentally, derive in part from the work of Alexander Graham Bell and that of his father A. Melville Bell. We turn, therefore, at this point to a discussion of the Bells' contribution to the science of language.

As is well known, Alexander Graham Bell was a speech therapist by profession. His specialty was the teaching of speech to the deaf, and according to all reports he was an extraordinarily gifted and successful practitioner of this difficult art. Speech therapy was the profession of many members of the Bell family, as shown on the bottom of the accompanying advertisement, which A. Melville Bell included at the end of his book *Visible Speech*. Speech therapy was a sort of family enterprise which the head of the family practiced in London and other members in other parts of Great Britain. What differentiated A. Melville Bell from most speech therapists was that he was interested not only in the practical aspects of his work but also in its scientific foundations. As we shall see, he involved his son in this work, the future inventor of the telephone, and on one issue of importance the latter made a contribution that went far beyond that of his father.

A. Melville Bell's analysis of spoken language proceeds from the observation that the production of speech sounds involves the coordinated activity of a number of different organs such as the lips, the tongue, the velum, and the larynx, which together make up what traditionally has been called the human vocal tract. From this point of view, the act of speaking is an elaborate gymnastics or choreography executed by different speech organs. In the book *Visible Speech*, we find a systematic account of the different activities that each speech organ is capable of, together with a discus-

sion of the different speech sounds that result from particular combinations of activities of specific speech organs.

Consider from this point of view the speech sounds that are produced by blowing air through a narrow opening as found in the words

veal zeal sheep wheel what

Sounds produced in this fashion are called *continuants*. One of the things that differentiate one continuant from another is the organ or organs actively involved in its formation, especially the constrictions—that is, the places in the vocal tract that are maximally narrowed when the sound in question is produced—and the organs actively effecting the narrowing, as shown

DIAGRAMS SHOWING THE RELATION OF THE PRIMARY ORGANIC SYMBOLS TO THE ORGANS.

CONSONANTS.

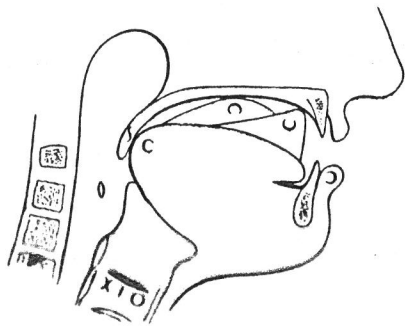

X Glottis closed, (catch.)
I " narrow, (voice.)
O " open, (aspirate.)
0 Super-Glottal Passage contracted, (whisper.)
(Soft Palate depressed, (nasal.)
C Back of Tongue, (contracting oral passage.)
C Front of do. (do.)
C Point of do. (do.)
C Lips, (do.)

in the drawing, which is reproduced from *Visible Speech*. Bell distinguished basically four constrictions: in /*f*/ * the constriction is formed by raising the lower lip; in /*z*/ and /*s*/ the constriction is formed by raising the blade of the tongue, whereas in /*ç*^{w}/ and /*x*^{w}/ there are two constrictions, one formed with the lower lip and the other with the tongue body, or *dorsum*. A further mechanism that is involved in distinguishing one sound from another is whether or not the sound is produced with the accompaniment of vocal cord vibration: /*z v*/ are; /*s ç*^{w} *x*^{w}/ are not. This fact can readily be verified by placing one's fingertips on the large (thyroid) cartilage in the front of the neck and pronouncing the sounds in question. When the vocal cords vibrate, this can be detected by a slight throbbing sensation in the fingertips. Finally, for purposes of this discussion, we need to identify one additional mechanism. It is the mechanism that produces strident sounds, such as /*f v s z č ǰ*/ and distinguishes them from the rest. It consists in directing the air stream against the sharp edges of the upper teeth, thereby producing audible turbulence.

We have thus identified five distinct mechanisms that are involved in the production of the continuant sounds under discussion. We label these for present purposes as follows:

- The raising of the lower lip—labial
- The raising of the tongue blade—coronal
- The raising of the tongue body—dorsal
- Vocal vibration—voicing
- Air stream directed at upper teeth—strident

When two or more mechanisms are activated the perceptual effect is that of a single sound. Thus, both /*z*/

* Slash marks mean the sound represented by the letter

as in *zeal* and /s/ as is *seal* are perceived as single sounds, although in the production of /z/ one more mechanism (voicing) is activated than in the production of /s/. As shown in the drawing, the *Visible Speech* alphabet had a special symbol to represent each of these mechanisms,—for example, the labial mechanism is represented there by a semi-circle open to the left, the coronal mechanism by a semi-circle open to the top, etc. When two or more mechanisms are activated in the production of a given sound the symbolic representation becomes rather cumbersome. It is therefore more convenient to represent the same information by means of a matrix such as the one in the accompanying chart (on page 51).

The claim made explicitly by A. Melville Bell in his *Visible Speech* is that he had identified all mechanisms that are relevant in the production of sounds in any spoken language. If this claim is correct, then it should be possible for an appropriately trained person to analyze any sound whatever in terms of the mechanisms involved in its production, especially since the number of mechanisms is fairly small. Moreover, it should also be possible for a trained person to produce sounds represented in this notation that he had never heard before. That is exactly how Bell saw the matter and he set about demonstrating it in a most dramatic fashion. I quote from *Visible Speech* (p. 22) :

> For the sake of showing the mode in which the experiments were conducted, the following description is quoted from a letter to the "Reader," by Alexander J. Ellis, Esq., F.R S.:—
>
> The mode of procedure was as follows: Mr. Bell sent his two Sons, who were to read the writing, out of the room,—it is interesting to know that the elder, who read all the words in this case, had only five

> weeks' instruction in the use of the Alphabet,—and I dictated slowly and distinctly the sounds which I

		LABIAL	CORONAL	DORSAL	VOICED	STRIDENT
f	*feel*	+	−	−	−	+
v	*veal*	+	−	−	+	+
x^w	*what*	+	−	+	−	−
s	*seal*	−	+	−	−	+
z	*zeal*	−	+	−	+	+
š	*she'll*	−	+	−	−	+
ž	*rouge*	−	+	−	+	+
č	*cheap*	−	+	−	−	+
ǰ	*jeep*	−	+	−	+	+
x	*Bach*	−	−	+	−	−
p	*peal*	+	−	−	−	−
d	*deal*	−	+	−	+	−
k	*keel*	−	−	+	−	−

wished to be written. These consisted of a few words in Latin, pronounced first as at Eton, then as in Italy, and then according to some theoretical notions of how Latins might have uttered them. Then came some English provincialisms and affected pronunciations; the words 'how odd' being given in several distinct ways. Suddenly German provincialisms were introduced. Then discriminations of sounds often confused. . . . Some Arabic, some Cockney-English, with an introduced Arabic guttural, some mispronounced Spanish, and a variety of vowels and diphthongs. . . . The result was perfectly satisfactory: —that is, Mr. Bell wrote down my queer and purposely-exaggerated pronunciations and mispronunciations, and delicate distinctions, in such a manner that his Sons, not having heard them, so uttered them as to surprise me by the extremely correct echo of my own voice. . . . Accent, tone drawl, brevity, indistinctness, were all reproduced with surprising accuracy. Being on the watch, I could, as it were, trace the alphabet in the lips of the readers. I think, then, that Mr. Bell is justified in the somewhat bold title which he has assumed for his mode of writing—"Visible Speech."

The quaintness of this testimonial should not be permitted to obscure the serious point that Bell attempted to establish by means of his demonstration, namely, that all sounds of all known languages can be produced given the very restricted information about a small number of mechanisms that is provided by *Visible Speech*. Anybody who controls all the mechanisms singly and in combination can produce any speech sound whatever. It is, therefore, these mechanisms and not the individual sounds of language that are the

fundamental building blocks of speech. This insight, which in the last quarter-century has become almost a truism among students of language, was stated explicitly in the early 1900's by Alexander Graham Bell in a series of lectures that he delivered to the American Association to Promote the Teaching of Speech to the Deaf. (It should be noted that Bell's terms *constriction* and *position* are synonymous with what has been termed *mechanism* here.)

> What we term an element of speech may in reality . . . be a combination of positions. The true element of articulation, I think, is a constriction or position of the vocal organs rather than a sound. Combinations of positions yield new sounds, just as combinations of chemical elements yield new substances. Water is a substance of very different character from either of the gases of which it is formed; and the vowel *oo* is a sound of very different character from that of any of its elementary positions.
>
> When we symbolize positions, the organic relations of speech sounds to one another can be shown by means of an equation; for example
> English $wh = P + P'$ *
> German $ch = P'$
> hence German ch = English $wh - P$
> The equation asserts that the English *wh* without labial constriction is the German *ch*. [*The Mechanism of Speech*, pp. 38–39]

I remarked above that during the last quarter-century it had become almost a truism among students of language that the elementary building blocks of lan-

* The symbol P in Bell's usage represents the phonetic feature labiality and P' represents the phonetic feature dorsality.

guage are not the sounds but the mechanisms or features as they are now called. While this idea has obvious plausibility as far as the physical production of speech is concerned, it is by no means self-evident that it is useful for other aspects of language. I shall not attempt to support this proposition directly. I shall deal with the data in one case as if the proposition were true and compare it to a solution where this assumption is not made. It will turn out that the obviously superior solution is that one couched in terms of features rather than in terms of speech sounds. I take this as partial evidence in support of the view that features rather than sounds are the basic elements of language in all its manifestations.

1. a) bus bush batch buzz garage judge
   b) cup cut cake cough sixth
   c) cab cad cog cam can song sea shoe sill flower

If you say to yourself the plural forms of the words in (1) you will notice that there is not one but three plural suffixes in English, one for each of the three separate sets of words in (1). We add an extra syllable */iz/* in forming the plural of the words in (1a) ; we add */s/* for the plural of the words in (1b), and we add */z/* to form the plural of the words in (1c). One can readily show that it is not the case that we memorize the plural form of every word we learn, for we know how to form the plurals of words we have never encountered before. Specifically, think of the plurals of the three English words mentioned earlier:

flitch plast thole

I am sure that most people here would agree that they know the plural forms of these previously unheard words and that these are respectively

flitches (like busses [1a])
plasts (like cups [1b])
tholes (like cabs [1c])

This forces us to conclude that speakers of English know a rule for the formation of plurals of nouns in spite of the fact that they have not been taught this rule by their parents.

It is necessary to be clear about the status of a rule such as the plural rule under discussion here. They are part of the knowledge that English speakers have and that people who do not know English normally do not possess. Knowing the rule that determines the phonetic actualization of the plural in English is, therefore, much like knowing that the device whose invention we commemorate in 1976 is called *telephone* rather than *farspeaker* (*cf. loudspeaker*), *phonex* or *glub*. The main difference between knowing the rule for the plural and knowing the word *telephone* is that the latter is conscious knowledge about which the speaker can answer direct questions, whereas knowledge of the plural rule and similar matters is largely unconscious and conceivably might never be accessible to consciousness. This fact, it should be noted at once, does not render such knowledge inaccessible to psychologists or linguists—that is, to scientists whose subject of inquiry is the speaker and his knowledge. Tacit knowledge can be established by the same methods that were used to establish other things inaccessible to direct observation such as the nature of the chemical bond or the structure of the gene.

The question that we want to answer is, In what form does the English speaker internalize his knowledge of the plural rule? An obvious candidate is:

2)
   a) If the noun ends with */s z š ž č ǰ/*, add */iz/*.

b) Otherwise, if the noun ends with /*p t k f θ*/, add /*s*/.
c) Otherwise, add /*z*/.

It is important to note about this rule that it is formulated in terms of speech sounds rather than in terms of mechanisms or features. In the light of the above discussion which suggested that features rather than sounds are the ultimate constituents of language, an attempt might be made to reformulate the rule in terms of features. The first move that one might make might be to replace each of the alphabetic symbols in (b) by its feature composition as shown in the preceding chart. Specifically, this means that one might replace /*s*/ by the feature complex /nonlabial-coronal-nondorsal-nonvoiced-strident/ /*z*/ by the same set of features except that in place of nonvoiced it would contain the feature voiced, etc. It is not easy to see where such a translation of the rule into feature terminology gets us. In fact, it gets us nowhere until we observe that with a chart such as that given earlier, it is possible to designate groups of sounds by mentioning one or two features. Thus, for example, if we asked for all and only sounds that are labial we would get the group /*f v x*^{w} *p*/; whereas if we asked for the sounds that are strident we would get /*f v s z s z č ǰ*/. Suppose now that we were to utilize this idea in the formulation of the plural rule and characterize each of the different lists of sounds by the minimum number of features that suffice to designate the group unambiguously. We should then get in place of (2),

3)
a) If the noun ends with a sound that is /coronal-strident/, add /iz/.

b) Otherwise, if the noun ends with a sound that is /nonvoiced/, add /*s*/.
c) Otherwise, add /*z*/.

Having formulated an alternative to the rule given above as (2), our task now is to determine which of the two alternatives is the one that English speakers use. The test that we shall use is one suggested to me some years ago by Lise Menn. It consists of asking English speakers to form the plural of a foreign word which ends with a sound that does not occur in English. A good example, Ms. Menn suggested, is the German name *Bach,* as in Johann Sebastian. If English speakers were operating in accordance with rule (2), they would have to reject options (a) and (b) and form the plural in accordance with option (c); that is, they would say that the plural of /*bax*/ is /*baxz*/ with a word final /*z*/. If, on the other hand, English speakers were operating in accordance with rule (3), they would have to perform a feature analysis of /*x*/ which would tell them that the sound is /nonlabial-noncoronal-dorsal-nonvoiced-nonstrident/. Given this feature composition, the plural of /*bax*/ could not be formed in accordance with option (a), since /*x*/ is neither coronal nor strident; it would, however, have to be formed in accordance with option (b) since /*x*/ is /nonvoiced/. In other words, if speakers operated in conformity with rule (3), their output would be /*baxs*/, which, as is perfectly obvious, is also the response that the majority of English speakers would make. We must, therefore, conclude that the formulation (3) of the plural rule in terms of features, and not the formulation (2) in terms of speech sounds, correctly represents the knowledge of English speakers.

There is yet another, more important inference to

be drawn from the fact that English speakers can apply the plural rule to a word ending with a sound that is not part of the repertory of English. In order to apply the rule, the speaker has to be able to establish that the foreign sound in question is nonvoiced. He must therefore have knowledge that allows him to determine the phonetic mechanism involved in the production of a sound that is not part of his language. The curious thing about such knowledge is that not only is there no indication that it might ever have been taught to speakers, there is also no indication that speakers could ever have acquired such knowledge. Think what evidence would have to be marshaled to support the claim that the knowledge in question was acquired. One would have to point to experiences in the life of the average English speaker that would permit him to acquire knowledge that is otherwise possessed only by phoneticians who have undergone rigorous training of the type Alexander Graham Bell received from his father. As this is obviously implausible, one is led to contemplate the possibility that at least some knowledge available to speakers is innate. In fact, there appears to be a certain amount of independent evidence that knowledge of the feature composition of sounds is available to children long before they could possibly have learned a language. Experiments conducted by Peter Eimas at Brown University and by Earl Butterfield at the University of Kansas have established that the ability to discriminate voiced from nonvoiced speech sounds is present in children practically at birth, but this ability presupposes knowledge which allows persons to determine the feature composition of speech sounds, that is, the same knowledge that is required in order to account for the ability of English speakers to form plurals of words with non-English sounds. The suggestion that this type of knowledge might be innate is, therefore, far from implausible.

This brings me to the end of what I have to say about the knowledge that speakers have of their language. What remains for me to do is to indicate how the information we have just reviewed helps us in trying to understand manifestations of the human cognitive capacity in domains other than language, how it might help us understand the human capacity to draw inferences, perform computations, play games with elaborate rules, interact with one another and uncover significant truths about the nature of the world around us and within us. If these manifestations of man's mind are at all like language, then we must expect to find that large portions of the knowledge on which they are based will be inaccessible to consciousness, that some of this knowledge will be innate, and that only a modest fraction of the total will have been acquired as the result of overt teaching. I must confess that I felt somewhat uneasy when I noticed that I was drawing attention to the possibility that teaching might play only a marginal role in the acquisition of knowledge. I was concerned about the effects that this remark might have on next year's enrollments in the M.I.T. linguistics program. But then I recalled the effect that the Surgeon General's warning printed on every pack of cigarettes has had on the popularity of smoking in this country and concluded that there was nothing to worry about since very little indeed is learned as the result of direct instruction.