In: J. F. Kavanaugh and Mattingly, I.G. (eds.), Language by Ear and by Eye. The Relationship between Speech and Reading (1972). MIT Press, Cambridge, MA and London, England.

IGNATIUS G. MATTINGLY

Reading, the Linguistic Process, and Linguistic Awareness

Reading is a rather remarkable phenomenon. The more we learn about speech and language, the more it appears that linguistic behavior is highly specific. The possible forms of natural language are very restricted; its acquisition and function are biologically determined [Chomsky 1965]. There is good reason to believe that special neural machinery is intricately linked to the vocal tract and the ear, the output and input devices used by all normal human beings for linguistic communication [Liberman, Cooper et al. 1967]. It is therefore rather surprising to find that a substantial number of human beings can also perform linguistic functions by means of the hand and the eye. If we had never observed actual reading or writing we would probably not believe these activities to be possible. Faced with the fact, we ought to suspect that some special kind of trick is involved. What I want to discuss is this trick, and what lies behind it—the relationship of the process of reading a language to the processes of speaking and listening to it. My view is that this relationship is much more devious than it is generally assumed to be. Speaking and listening are primary linguistic activities; reading is a secondary and rather special sort of activity that relies critically upon the reader's awareness of these primary activities.

Comparison of Reading and Listening

The usual view, however, is that reading and listening are parallel processes. According to this view, written text is input by eye and speech by ear, but at as early a stage as possible, consistent with this difference in modality, the two inputs have a common internal representation. From this stage onward, the two processes are identical. Reading is ordinarily learned later than speech; this learning is therefore essentially an intermodal transfer, the attainment of skill in doing visually what one already knows how to do auditorily. As C. C. Fries [1962, p. xv] puts it, "Learning to read . . . is *not* a process of learning new or other language signals than those the child has already learned. The language signals are all the same. The difference lies in the medium through which the physical stimuli make contact with his nervous system. In 'talk' the physical stimuli of the language signals make their contact by means of sound waves received by the ear. In reading, the physical

stimuli of the same language signals consist of graphic shapes that make their contact with his nervous system through light waves received by the eye. The process of learning to read is the process of transfer from the auditory signs for language signals which the child has already learned, to the new visual signs for the same signals."

Something like this view appears to be shared by many who differ about other aspects of reading, even about the nature of the linguistic activity involved. Thus Bloomfield [1955], Fries and others assume that the production and perception of speech are inversely related processes of encoding and decoding, and take the same view of writing and reading. They believe that the listener extracts the phonemes or "unit speech sounds" from speech, forms them into morphemes and sentences, and decodes the message. Similarly, the reader produces, in response to the text, either audible unit speech sounds or, in silent reading, "internal substitute movements" [Bloomfield 1955, p. 103] which he treats as phonemes and so decodes the message. Fries's model is similar to Bloomfield's except that his notion of a phoneme is rather more abstract; it is a member of a set of contrasting elements, conceptually distinct from the medium which conveys it. This medium is the acoustic signal for the listener, the line of print for the reader. For Fries as for Bloomfield, acquisition of both the spoken and written language requires development of "high-speed recognition responses" to stimuli which "sink below the threshold of attention" [Fries 1962, p. xvi) when the responses have become habitual.

More recently, however, the perception of speech has come to be regarded by many as an "active" process basically similar to speech production. The listener understands what is said through a process of "analysis by synthesis" [Stevens and Halle 1967]. Parallel proposals have accordingly been made for reading. Thus Hochberg and Brooks [1970] suggest that once the reader can visually discriminate letters and letter groups and has mastered the phoneme-grapheme correspondences of his writing system, he uses the same hypothesis-testing procedure in reading as he does in listening (Goodman's [1970] view of reading as a "psycholinguistic guessing game" is a similar proposal). Though the model of linguistic processing is different from that of Bloomfield and Fries, the assumption of a simple parallel between reading and listening remains, and the only differences mentioned are those assignable to modality, for example, the use which the reader makes of peripheral vision, which has no analog in listening.

While it is clear that reading somehow employs the same linguistic processes as listening, it does not follow that the two activities are directly

analogous. There are, in fact, certain differences between the two processes that cannot be attributed simply to the difference of modality, and which therefore make difficulties for the notion of a straightforward intermodal parallel. Most of these differences have been pointed out before, notably by Liberman, Cooper, Shankweiler and Studdert-Kennedy [1967] and by Liberman at an earlier conference on "the Reading Process" [Kavanagh 1968]. But I think reconsideration of them will help us to arrive at a better understanding of reading.

To begin with, listening appears to be a more natural way of perceiving language than reading; "listening is easy and reading is hard" [Liberman, in Kavanagh 1968, p. 119]. We know that all living languages are spoken languages, and that every normal child gains the ability to understand his native speech as part of a maturational process of language acquisition. In fact we must suppose that, as a prerequisite for language acquisition, the child has some kind of innate capability to perceive speech. In order to extract from the utterances of others the "primary linguistic data" that he needs for acquisition, he must have a "technique for representing input signals" [Chomsky 1965, p. 30].

In contrast, relatively few languages are written languages. In general, children must be deliberately taught to read and write, and despite this teaching, many of them fail to learn. Someone who has been unable to acquire language by listening—a congenitally deaf child, for instance—will hardly be able to acquire it through reading; on the contrary, as Liberman and Furth [Kavanagh 1968] point out, a child with a language deficit owing to deafness will have great difficulty learning to read properly.

The apparent naturalness of listening does not mean that it is in all respects a more efficient process. Though many people find reading difficult, there are a few readers who are very proficient: in fact, they read at rates well over 2000 words per minute with complete comprehension. Listening is always a slower process: even when speech is artificially speeded up in a way which preserves frequency relationships, 400 words per minute is about the maximum possible rate [Orr, Friedman et al. 1965]. It has often been suggested [e.g., Bever and Bower 1966; Bower, 1970] that high-speed readers are somehow able to go directly to a deep level of language, omitting the intermediate stages of processing to which other readers and all listeners must presumably have recourse.

Moreover, the form in which information is presented is basically different in reading and in listening. The listener is processing a complex acoustic signal in which the speech cues that constitute significant linguistic data are buried. Before he can use these cues, the listener has

to "demodulate" the signal: that is, he has to separate the cues from the irrelevant detail. The complexity of this task is indicated by the fact that no scheme for speech recognition by machine has yet been devised that can perform it properly. The demodulation is largely unconscious; as a rule, a listener is unable to perceive the actual acoustic form of the event which serves as a cue unless it is artificially excised from its speech context [Mattingly, Liberman et al. 1971]. The cues are not discrete events well separated in time or frequency; they blend into one another; we cannot, for instance, realistically identify a certain instant as the ending of a formant transition for an initial consonant and the beginning of the steady state of the following vowel (see Cooper's paper, this volume).

The reader, on the other hand, is processing a series of symbols that are quite simply related to the physical medium that conveys them. The task of demodulation is straightforward: the marks in black ink are information; the white paper is background. The reader has no particular difficulty in seeing the letters as visual shapes if he wants to. In printed text, the symbols are discrete units. In cursive writing, of course, one can slur together the symbols to a surprising degree without loss of legibility. But though they are deformed, the cursive symbols remain essentially discrete. It makes sense to view cursive writing as a string of separate symbols connected together for practical convenience; it makes no sense at all to view the speech signal in this way.

That these differences in form are important is indicated by the difficulty of reading a visual display of the speech signal, such as a sound spectrogram, or of listening to text coded in an acoustic alphabet, for example, Morse code or any of the various acoustic alphabets designed to aid the blind [Studdert-Kennedy and Liberman 1963; Coffey 1963]. We know that a spectrogram contains most of the essential linguistic information, for it can be converted back to acoustic form without much loss of intelligibility [Cooper 1950]. Yet reading a spectrogram is very slow work at best, and at worst impossible. Similarly, text coded in an acoustic alphabet contains the same information as print, but a listener can decode it only if it is presented at a rate which is very slow compared to a normal speaking rate.

These facts are certainly not quite what we should predict if reading and listening were simply similar processes in different modalities. The relative advantage of the eye with alphabetic text, to be sure, may be attributed to its apparent superiority over the ear as a data channel; but then why should the eye do so poorly with visible speech? We can

only infer that some part of the neural speech processing machinery must be accessible through the ear but not through the eye.

There is also a difference in the linguistic content of the information available to the listener and the reader. The speech cues carry information about the phonetic level of language, the articulatory gestures which the speaker must have made—or more precisely, the motor commands which lead to those gestures [Lisker, Cooper et al. 1962]. Written text corresponds to a different level of language. Chomsky [1970] makes the important observation that conventional orthography, that of English in particular, is, roughly speaking, a morphophonemic transcription; in the framework of generative grammar, it corresponds fairly closely to a surface-structure phonological representation; Chomsky uses the term "lexical representation." I think this generalization can probably be extended to include all practical writing systems, despite their apparent variety. The phonological level is quite distinct from the phonetic level though the two are linked in each language by a system of phonological rules. The parallel between listening and reading was plausible in part because of the failure of structural linguistics to treat these two linguistic levels as the significant ones: both speech perception and reading were taken to be phonemic. Chomsky [1964] and Halle [1959], however, have argued rather convincingly that the phonemic level of the structuralists has no proper linguistic significance, its supposed functions being performed either at the phonological or the phonetic levels.

Halwes observed at the conference on the Reading Process [Kavanagh 1968, p. 160] "It seems like a good bet that since you have all this apparatus in the head for understanding language that if you wanted to teach somebody to read, you would arrange a way to get the written material input to the system that you have already got for processing spoken language and at as low a level as you could arrange to do that, then let the processing of the written material be done by the mechanisms that are already in there." I think that Halwes' inference is a reasonable one, and since the written text is not, in fact, a representation at the lowest possible level, the problem is with his premise, that reading and listening are simply analogous processes.

There is furthermore a difference in the way the linguistic content and the information which represents it are related. As Liberman [Kavanagh 1968, p. 120] observes, "speech is a complex code, print a simple cipher." The nature of the speech code by which the listener deduces articulatory behavior from acoustic events is determined by the characteristics of the vocal tract. The code is complex, because the physi-

ology and acoustics of the vocal tract are complex. It is also a highly redundant code: there are, typically, many acoustic cues for a single bit of phonetic information. It is, finally, a universal code, because all human vocal tracts have similar properties. By comparison, writing is in principle a fairly simple mapping of units of the phonological representation—morphemes or phonemes or syllables—into written symbols. The complications that do occur are not inherent in the nature of what is being represented: they are historical accidents. By comparison with the speech code, writing is a very economical mapping; typically, many bits of phonological information are carried by a single symbol. Nor is there any necessary relationship between the form of written symbols and the corresponding phonological units; to quote Liberman once more [Kavanagh 1968, p. 121], "only one set of sounds will work, but there are many equally good alphabets."

Linguistic Awareness

The differences we have listed indicate that even though reading and listening are both clearly linguistic, and have an obvious similarity of function, they are not really parallel processes. I would like to suggest a rather different interpretation of the relationship of reading to language. This interpretation depends on a distinction between primary linguistic activity itself and the speaker-hearer's awareness of this activity.

Following Miller and Chomsky [1963], Stevens and Halle [1967], Neisser [1967], and others, I view primary linguistic activity, both speaking and listening, as essentially creative or synthetic. When a speaker-hearer "synthesizes" a sentence, the products are a semantic representation and a phonetic representation that are related by the grammatical rules of his language, in the sense that the generation of one entails the generation of the other. The speaker must synthesize, and so produce a phonetic representation for, a sentence which, according to the rules, will have a particular required semantic representation; the listener, similarly, must synthesize a sentence which matches a particular phonetic representation, in the process recovering its semantic representation. It should be added that synthesis of a sentence does not necessarily involve its utterance. One can think of a sentence without actually speaking it; one can rehearse or recall a sentence.

Since we are concerned with reading and not with primary linguistic activity as such, we will not attempt the difficult task of specifying the actual process of synthesis. We merely assume that the speaker-hearer not only knows the rules of his language but has a set of strategies for linguistic performance. These strategies, relying upon context as well

as upon information about the phonetic (or semantic) representation to be matched, are powerful enough to ensure that the speaker-hearer synthesizes the "right" sentence most of the time.

Having synthesized some utterance, whether in the course of production or perception, the speaker-hearer is conscious not only of a semantic experience (understanding the utterance) and perhaps an acoustic experience (hearing the speaker's voice), but also of experience with certain intermediate linguistic processes. Not only has he synthesized a particular utterance, he is also aware in some way of having done so, and can reflect upon this linguistic experience as he can upon his experiences with the external world.

If language were in great part deliberately and consciously learned behavior, like playing the piano, this would hardly be very surprising. We would suppose that development of such linguistic awareness was needed in order to learn language. But if language is acquired by maturation, linguistic awareness seems quite remarkable when we consider how little introspective awareness we have of the intermediate stages of other forms of maturationally acquired motor and perceptual behavior, for example, walking or seeing. (The concept of "linguistic awareness" developed here is similar but not identical to the concept of "accessibility" discussed in Klima's paper in this volume.)

The speaker-hearer's linguistic awareness is what gives linguistics its special advantage in comparison with other forms of psychological investigation. Taking his informant's awareness of particular utterances as a point of departure, the linguist can construct a description of the informant's intuitive competence in his language which would be unattainable by purely behavioristic methods [Sapir 1949].

However, linguistic awareness is very far from being evenly distributed over all phases of linguistic activity. Much of the process of synthesis takes place well beyond the range of immediate awareness [Chomsky 1965], and must be determined inferentially—just how much has become clear only recently, as a result of investigations of deep syntactic structure by generative grammarians and of speech perception by experimental phoneticians. Thus the speaker-hearer's knowledge of the deep structure and transformational history of an utterance is evident chiefly from his awareness of the grammaticality of the utterance or its lack of it; and he has no direct awareness at all of many of the most significant acoustic cues, which have been isolated by means of perceptual experiments with synthetic speech.

On the other hand, the speaker-hearer has a much greater awareness of phonetic and phonological events. At the phonetic level, he can often

detect deviations, even in the case of features which are not distinctive in his language, and this sort of awareness can be rapidly increased by appropriate ear training.

At the phonological (surface-structure) level, not only distinctions between deviant and acceptable utterances, but also reference to various structural units become possible. Words are perhaps most obvious to the speaker-hearer, and morphemes hardly less so, at least in the case of languages with fairly elaborate inflectional and compounding systems. Syllables, depending upon their structural role in the language, may be more obvious than phonological segments. There is far greater awareness of the structural unit than of the structure itself, so that the speaker-hearer feels that the units are simply concatenated. The syntactic bracketing of the phonological representation is probably least obvious.

In the absence of appropriate psycholinguistic data (see, however, Savin's paper in this volume), any ordering of this sort is of course very tentative, and in any case, it would be a mistake to overstate the clarity of the speaker-hearer's linguistic awareness and the consistency with which it corresponds to a particular linguistic level. But it is safe to say that, by virtue of this awareness, he has an internal image of the utterance, and this image probably owes more to the phonological level of representation than to any other level.

There appears to be considerable individual variation in linguistic awareness. Some speaker-hearers are not only very conscious of linguistic patterns but exploit their consciousness with obvious pleasure in verbal play (e.g., punning) or verbal work (e.g., linguistic analysis). Others seem never to be aware of much more than words and are surprised when quite obvious linguistic patterns are pointed out to them. This variation contrasts markedly with the relative consistency from person to person with which primary linguistic activity is performed. Synthesis of an utterance is one thing; the awareness of the process of synthesis quite another.

Linguistic awareness is by no means only a passive phenomenon. The speaker-hearer can use his awareness to control, quite consciously, his linguistic activity. Thus he can ask himself to synthesize a number of words containing a certain morpheme, or a sentence in which the same phonological segment recurs repeatedly. Without this active aspect of linguistic awareness, moreover, much of what we call thinking would be impossible. The speaker-hearer can consciously represent things by names and complex concepts by verbal formulas. When he tries to think abstractly, manipulating these names and concepts, he relies ultimately upon his ability to recapture the original semantic experiences they repre-

sent. The only way to do this is to resynthesize the utterance to which a name or formula corresponds.

Moreover, linguistic awareness can become the basis of various language-based skills. Secret languages, such as Pig Latin [Halle 1964] form one class of examples. In such languages a further constraint, in the form of a rule relating to the phonological representation, is artificially imposed upon production and perception. Having synthesized a sentence in English, an additional mental operation is required to perform the encipherment. To carry out the process at a normal speaking rate, one has not only to know the rule but also to have developed a certain facility in applying it. A second class of examples are the various systems of versification. The versifier is skilled in synthesizing sentences which conform not only to the rules of the language but to an additional set of rules relating to certain phonetic features [Halle 1970]. To listen to verse, one needs at least a passive form of this skill so that one can readily distinguish "correct" from "incorrect" lines without scanning them syllable by syllable.

It seems to me that there is a clear difference between Pig Latin, versification, and other instances of language-based skills, and primary linguistic activity itself. If one were unfamiliar with Pig Latin or with a system of versification, one might fail to understand what the Pig Latinist or the versifier was up to, but one would not suppose either of them to be speaking an unfamiliar language. And even after one does get on to the trick, the sensation of engaging in something beyond primary linguistic activity does not disappear. One continues to be aware of a special demand upon one's linguistic awareness.

Reading as a Language-Based Skill

Our view is that reading is a language-based skill like Pig Latin or versification and not a form of primary linguistic activity analogous to listening. From this viewpoint, let us try to give an account, necessarily much oversimplified, of the process of reading a sentence.

The reader first forms a preliminary, quasiphonological representation of the sentence based on his visual perception of the written text. The form in which this text presents itself is determined not by the actual linguistic information conveyed by the sentence but by the writer's linguistic awareness of the process of synthesizing the sentence, an awareness which the writer wishes to impart to the reader. The form of the text does *not* consist, for instance, of a tree-structure diagram or a representation of articulatory gestures, but of discrete units, clearly separable from their visual context. These units, moreover, correspond roughly to ele-

ments of the phonological representation (in the generative grammarian's sense), and the correspondence between these units and the phonological elements is quite simple. The only real question is whether the writing system being used is such that the units represent morphemes, or syllables, or phonological segments.

Though the text is in a form which appeals to his linguistic awareness, considerable skill is required of the reader. If he is to proceed through the text at a practical pace, he cannot proceed unit by unit. He must have an extensive vocabulary of sight words and phrases acquired through previous reading experience. Most of the time he identifies long strings of units. When this sight vocabulary does fail him, he must be ready with strategies by means of which he can identify a word that is part of his spoken vocabulary and add it to his sight vocabulary, or assign a phonological representation to a word altogether unknown to him. To be able to do this he must be thoroughly familiar with the rules of the writing system: the shapes of the characters and the relationship of characters and combinations of characters to the phonology of his language. Both sight words and writing system are matters of convention, and must be more or less deliberately learned. While their use becomes habitual in the skilled reader, they are never inaccessible to awareness in the way that much primary linguistic activity is.

The preliminary representation of the sentence will contain only a part of the information in the linguist's phonological representation. All writing systems omit syntactic, prosodic, and junctural information, and many systems make other omissions: for example, phonological vowels are inadequately represented in English spelling and omitted completely in some forms of Semitic writing. Thus the preliminary representation recovered by the reader from the written text is a partial version of the phonological representation: a string of words which may well be incomplete and are certainly not syntactically related.

The skilled reader, however, does not need complete phonological information, and probably does not use all of the limited information available to him. The reason is that the preliminary phonological representation serves only to control the next step of the operation, the actual synthesis of the sentence. By means of the same primary linguistic competence he uses in speaking and listening, the reader endeavors to produce a sentence that will be consistent with its context and with this preliminary representation. In order to do this, he needs, not complete phonological information, but only enough to exclude all other sentences which would fit the context. As he synthesizes the sentence, the reader

derives the appropriate semantic representation, and so understands what the writer is trying to say.

Does the reader also form a phonetic representation? Though it might seem needless to do so in silent reading, I think he does. In view of the complex interaction between levels which must take place in primary linguistic activity, it seems unlikely that a reader could omit this step at will. Moreover, as suggested earlier, even though writing systems are essentially phonological, linguistic awareness is in part phonetic. Thus, a sentence that is phonetically bizarre—"The rain in Spain falls mainly in the plain," for example—will be spotted by the reader. And quite often, the reason a written sentence appears to be stylistically offensive is that it would be difficult to speak or listen to.

Having synthesized a sentence that fits the preliminary phonological representation, the reader proceeds to the actual recognition of the written text; that is, he applies the rules of the writing system and verifies, at least in part, the sentence he has synthesized. Thus we can, if we choose, think of the reading process as one analysis-by-synthesis loop inside another, the inner loop corresponding to primary linguistic activity and the outer loop to the additional skilled behavior used in reading. This is a dangerous analogy, however, because the nature of both the analysis and the synthesis is very different in the two processes.

This account of reading ties together many of the differences between reading and listening noted earlier: the differences in the form of the input information, the difference in its linguistic content, and the difference in the relationship of form to content. But we have still to explain the two most interesting differences: the relatively higher speeds that can be attained in reading and the relative difficulty of reading.

How can we explain the very high speeds at which some people read? To say that such readers go directly to a semantic representation, omitting most of the process of linguistic synthesis, is to hypothesize a special type of reader who differs from other readers in the nature of his primary linguistic activity, and differs in a way which we have no other grounds for supposing possible. As far as I know, no one has suggested that high-speed readers can *listen*, rapidly or slowly, in the way they are presumed to read. A more plausible explanation is that linguistic synthesis takes place much faster than has been supposed, and that the rapid reader has learned how to take advantage of this. The relevant experiments (summarized by Neisser [1967]) have measured the rate at which rapidly articulated or artificially speeded speech can be comprehended, and the rate at which a subject can count silently, that is, the rate

of "inner speech." But since temporal relationships in speech can only withstand so much distortion, speeded speech experiments may merely reflect limitations on the rate of input. The counting experiment not only used unrealistic material but assumed that inner speech is an essential concomitant of linguistic synthesis.

But suppose that the inner speech which so many readers report, and which figures so prominently in the literature on reading, is simply a kind of auditory imagery, dependent upon linguistic awareness of the sentence already synthesized, reassuring but by no means essential to synthesis (any more than actual utterance or subvocalization), and rather time-consuming. One could then explain the high speed reader as one who builds up the preliminary representation efficiently and synthesizes at a very high speed, just as any other reader or speaker-hearer does. But since he is familiar with the nature of the text, he seldom finds it necessary to verify the output of the process of synthesis, and spends no time on inner speech. The high speed at which linguistic synthesis occurs is thus directly reflected in his reading speed. This explanation is admittedly speculative but has the attraction of treating the primary linguistic behavior of all readers as similar, and assigning the difference to behavior peculiar to reading.

Finally, why should reading be, by comparison with listening, so perilous a process? This is not the place to attempt an analysis of the causes of dyslexia, but if our view of reading is correct, there is plenty of reason why things should often go wrong. First, we have suggested that reading depends ultimately on linguistic awareness and that the degree of this awareness varies considerably from person to person. While reading does not make as great a demand upon linguistic awareness as, say, solving British crossword puzzles, there must be a minimum level required, and perhaps not everyone possesses this minimum; not everyone is sufficiently aware of units in the phonological representation or can acquire this awareness by being taught. In the special case of alphabetic writing, it would seem that the price of greater efficiency in learning is a required degree of awareness higher than for logographic and syllabary systems, since as we have seen, phonological segments are less obvious units than morphemes or syllables. Almost any Chinese with ten years to spare can learn to read, but there are relatively few such people. In a society where alphabetic writing is used, we should expect more reading successes, because the learning time is far shorter, but proportionately more failures, too, because of the greater demand upon linguistic awareness.

A further source of reading difficulty is that the written text is a

grosser and far less redundant representation than speech: one symbol stands for a great deal more information than one speech cue, and the same information is not available elsewhere in the text. Both speaker and listener can perform sloppily and the message will get through: the listener who misinterprets a single speech cue will often be rescued by several others. Even a listener with some perceptual difficulty can muddle along. The reader's tolerance of noisy input is bound to be much lower than the listener's, and a person with difficulty in visual perception so mild as not to interfere with most other tasks may well have serious problems in reading.

These problems are both short-term and long-term. Not only does the poor reader risk misreading the current sentence, but there is the possibility that his vocabulary of sight words and phrases will become corrupted by bad data, and that the strategies he applies when the sight vocabulary fails will be the wrong strategies. In this situation he will build up the preliminary phonological representation not only inaccurately, which in itself might not be so serious, but also too slowly, because he is forced to have recourse to his strategies so much of the time. This is fatal, because a certain minimum rate of input seems to be required for linguistic synthesis. We know, from experience with speech slowed by inclusion of a pause after each word, that even when individual words are completely intelligible, it is hard for the listener to put the whole sentence together. If only a reader can maintain the required minimum rate of input, many of his perceptual errors can be smoothed over in synthesis; it is no doubt for this reason that most readers manage as well as they do. But if he goes too slowly, he may well be unable to keep up with his own processes of linguistic synthesis and will be unable to make sense at all out of what he reads.

Liberman has remarked that reading is parasitic on language [Kavanagh 1968]. What I have tried to do here, essentially, is to elaborate upon that notion. Reading is seen not as a parallel activity in the visual mode to speech perception in the auditory mode; there are differences between the two activities that cannot be explained in terms of the difference of modality. They can be explained only if we regard reading as a deliberately acquired, language-based skill, dependent upon the speaker-hearer's awareness of certain aspects of primary linguistic activity. By virtue of this linguistic awareness, written text initiates the synthetic linguistic process common to both reading and speech, enabling the reader to get the writer's message and so to recognize what has been written.

References

Bever, T. G., and T. G. Bower, 1966. How to read without listening. *Project Literacy Reports No. 6,* pp. 13–25.

Bloomfield, Leonard, 1955. Linguistics and reading. *Language Learning* 5:94–107.

Bower, T. G., 1970. Reading by eye. In *Basic Studies on Readings,* H. Levin and J. P. Williams (eds.), New York: Basic Books.

Chomsky, N., 1964. *Current Issues in Linguistic Theory.* The Hague: Mouton.

———, 1965. *Aspects of the Theory of Syntax.* Cambridge, Mass.: M.I.T. Press.

———, 1970. Phonology and reading. In *Basic Studies on Reading,* H. Levin and J. P. Williams (eds.), New York: Basic Books.

Coffey, J. L., 1963. The development and evaluation of the Battelle Aural Reading Device. *Proceedings of the International Congress of Technology and Blindness,* New York: American Foundation for the Blind.

Cooper, F. S., 1950. Spectrum analysis. *J. Acoust. Soc. Amer.* 22:761–762.

Fries, C. C., 1962. *Linguistics and Reading.* New York: Holt, Rinehart and Winston.

Goodman, K. S., 1970. Reading: a psycholinguistic guessing game. In *Theoretical Models and Processes of Reading.* H. Singer and R. B. Ruddell (eds.), Newark, Del.: International Reading Association.

Halle, M., 1959. *The Sound Pattern of Russian.* The Hague: Mouton.

———, 1964. On the bases of phonology. In *The Structure of Language,* J. A. Foder and J. J. Katz (eds.), Englewood Cliffs, N.J.: Prentice-Hall.

———, 1970. On metre and prosody. In *Progress in Linguistics,* M. Bierwisch and K. Heidolph (eds.), The Hague: Mouton.

Hochberg, J. and V. Brooks, 1970. Reading as an intentional behavior. In *Theoretical Models and Processes of Reading.* H. Singer and R. B. Ruddell (eds.), Newark, Del.: International Reading Association.

Kavanagh, J. F. (ed.), 1968. *Communicating by Language: The Reading Process.* Bethesda, Md.: National Institute of Child Health and Human Development.

Liberman, A. M., F. S. Cooper, D. P. Shankweiler, and M. Studdert-Kennedy, 1967. Perception of the speech code. *Psych. Rev.* 74:431–461.

Lisker, L., F. S. Cooper, and A. M. Liberman, 1962. The uses of experiment in language description. *Word* 18:82–106.

Mattingly, I. G., A. M. Liberman, A. K. Syrdal, and T. Halwes, 1971. Discrimination in speech and non-speech modes. *Cog. Psych.* 2:131–157.

Miller, G., and N. Chomsky, 1963. Finitary models of language users. In *Handbook of Mathematical Psychology,* R. D., Luce, R. R. Bush, and E. Galanter (eds.), New York: Wiley.

Neisser, U., 1967. *Cognitive Psychology.* New York: Appleton-Century-Crofts.

Orr, D. B., H. L. Friedman, and J. C. C. Williams, 1965. Trainability of listening comprehension of speeded discourse. *J. Ed. Psych.* 56:148–156.

Sapir, E., 1949. The psychological reality of phonemes. In *Selected Writings of Edward Sapir in Language, Culture, and Personality,* D. G. Mandelbaum (ed.), Berkeley: University of California Press.

Studdert-Kennedy, M., and A. M. Liberman, 1963. Psychological considerations in the design of auditory displays for reading machines. *Proceedings of the International Congress of Technology and Blindness,* New York: American Foundation for the Blind.

Stevens, K. N., and M. Halle, 1967. Remarks on analysis by synthesis and distinctive features. In *Models for the Perception of Speech and Visual Form,* W. Wathen-Dunn (ed.), Cambridge, Mass: M.I.T. Press.