

The Psycholinguists

On the New Scientists of Language — By GEORGE A. MILLER

PSYCHOLOGISTS have long recognised that human minds feed on linguistic symbols. Linguists have always admitted that some kind of psycho-social motor must move the machinery of grammar and lexicon. Sooner or later they were certain to examine their intersection self-consciously. Perhaps it was also inevitable that the result would be called “psycholinguistics.”

In fact, although the enterprise itself has slowly been gathering strength at least since the invention of the telephone, the name, in its unhyphenated form, is only about ten years old. Few seem pleased with the term, but the field has grown so rapidly and stirred so much interest in recent years that some way of referring to it is urgently needed. *Psycholinguistics* is as descriptive a term as any, and shorter than most.

Among psychologists it was principally the behaviourists who wished to take a closer look at language. Behaviourists generally try to replace anything subjective by its most tangible, physical manifestation, so they have had a long tradition of confusing thought with speech—or with “verbal behaviour,” as many prefer to call it. Among linguists it was principally those with an anthropological sideline who were most willing to collaborate, perhaps because as anthropologists they were sensitive to all those social and psychological processes that support our

linguistic practices. By working together they managed to call attention to an important field of scientific research and to integrate it, or at least to acquaint its various parts with one another, under this new rubric.¹

Interest in psycholinguistics, however, is not confined to psychologists and linguists. Many people have been stirred by splendid visions of its practical possibilities. One thinks of medical applications to the diagnosis and treatment of a heterogeneous variety of language disorders ranging from simple stammering to the overwhelming complexities of aphasia.² One thinks too of pedagogical applications, of potential improvements in our methods for teaching reading and writing, or for teaching second languages. If psycholinguistic principles were made sufficiently explicit, they could be imparted to those technological miracles of the twentieth century, the computing machines, which would bring into view a whole spectrum of cybernetic possibilities.³ We could exploit our electrical channels for voice communications more efficiently. We might improve and automate our dictionaries, using them for mechanical translation from one language to another. Perhaps computers could print what we say, or even say what we print, thus making speech visible for the deaf and printing audible for the blind. We might, in short, learn to adapt computers to dozens of our human purposes if only they could interpret our languages. Little wonder that assorted physicians, educators, philosophers, logicians, and engineers have been intrigued by this new adventure.

Of course, the realisation of practical benefits must await the success of the scientific effort; there is some danger that enthusiasm may colour our estimate of what can be accomplished. Not a few sceptics remain unconvinced; some can even be found who argue that success is impossible in principle. “Science,” they say, “can go only so far. . . .”

¹ A representative sample of research papers in this field can be found in *Psycholinguistics, a Book of Readings*, edited by S. Saporta (Holt, Rinehart & Winston, New York, 1962). R. Brown provides a readable survey from a psychologist's point of view in *Words and Things* (Free Press, Glencoe, Illinois, 1957).

² The CIBA Foundation Symposium, *Disorders of Language* (J. & A. Churchill, London, 1964) provides an excellent sample of the current status of medical psycholinguistics.

³ *Natural Language and the Computer*, edited by P. L. Garvin (McGraw-Hill, New York, 1963).

THE INTEGRATION of psycholinguistic studies has occurred so recently that there is still some confusion concerning its scope and purpose; efforts to clarify it necessarily have something of the character of personal opinion.⁴ In my own version, the central task of this new science is to describe the psychological processes that go on when people use sentences. The real crux of the psycholinguistic problem does not appear until one tries to deal with sentences, for only then does the importance of productivity become completely obvious. It is true that productivity can also appear with individual words, but there it is not overwhelming. With sentences, productivity is literally unlimited.

Before considering this somewhat technical problem, however, it might be well to illustrate the variety of processes that psycholinguists hope to explain. This can best be done if we ask what a listener can do about a spoken utterance, and consider his alternatives in order from the superficial to the inscrutable.

The simplest thing one can do in the presence of a spoken utterance is to listen. Even if the language is incomprehensible, one can still *hear* an utterance as an auditory stimulus and respond to it in terms of some discriminative set: how loud, how fast, how long, from which direction, etc.

Given that an utterance is heard, the next level involves *matching* it as a phonemic pattern in terms of phonological skills acquired as a user of the language. The ability to match an input can be tested in psychological experiments by asking listeners to echo what they hear; a wide variety of experimental situations—experiments on the perception of speech and on the note memorisation of verbal materials—can be summarised as tests of a person's ability to repeat the speech he hears under various conditions of audibility or delay.

If a listener can hear and match an utterance, the next question to ask is whether he will

accept it as a sentence in terms of his knowledge of grammar. At this level we encounter processes difficult to study experimentally, and one is forced to rely most heavily on linguistic analyses of the structure of sentences. Some experiments are possible, however, for we can measure how much a listener's ability to accept the utterance as a sentence facilitates his ability to hear and match it; grammatical sentences are much easier to hear, utter or remember than are ungrammatical strings of words, and even nonsense (*pirot, karol, elat*, etc.) is easier to deal with if it looks grammatical (*pirots karolise elatically*, etc.).⁵ Needless to say, the grammatical knowledge we wish to study does not concern those explicit rules drilled into us by teachers of traditional grammar, but rather the implicit generative knowledge that we all must acquire in order to use a language appropriately.

Beyond grammatical acceptance comes semantic interpretation: we can ask how listeners *interpret* an utterance as meaningful in terms of their semantic system. Interpretation is not merely a matter of assigning meanings to individual words; we must also consider how these component meanings combine in grammatical sentences. Compare the sentences: *Healthy young babies sleep soundly* and *Colourless green ideas sleep furiously*. Although they are syntactically similar, the second is far harder to perceive and remember correctly—because it cannot be interpreted by the usual semantic rules for combining the senses of adjacent English words.⁶ The interpretation of each word is affected by the company it keeps; a central problem is to systematise the interactions of words and phrases with their linguistic contexts. The lexicographer makes his major contribution at this point, but psychological studies of our ability to paraphrase an utterance also have their place.

At the next level it seems essential to make some distinction between interpreting an utterance and understanding it, for understanding frequently goes well beyond the linguistic context provided by the utterance itself. A husband greeted at the door by "I bought some electric light bulbs to-day" must do more than interpret its literal reference; he must understand that he should go to the kitchen and replace that burned-out lamp. Such contextual information lies well outside any grammar or lexicon. The listener can *understand* the function of an utterance in terms of contextual knowledge of the most diverse sort.

Finally, at a level now almost invisible through the clouds, a listener may *believe* that an utterance is valid in terms of its relevance to his own conduct. The child who says "I saw five lions in the garden" may be heard, matched,

⁴ My own opinions have been strongly influenced by Noam Chomsky. A rather technical exposition of this work can be found in Chapters 11-13 of the second volume of the *Handbook of Mathematical Psychology*, edited by R. D. Luce, R. R. Bush, and E. Galanter (Wiley, New York, 1963), from which many of the ideas discussed here have been drawn.

⁵ W. Epstein, "The Influence of Syntactical Structure on Learning," *American Journal of Psychology* (1961), vol. 74, pp. 80-85.

⁶ G. A. Miller and S. Isard, "Some Perceptual Consequences of Linguistic Rules," *Journal of Verbal Learning and Verbal Behaviour* (1963), vol. 2, pp. 217-228. J. J. Katz and J. A. Fodor have recently contributed a thoughtful discussion of "The Structure of Semantic Theory," *Language* (1963), vol. 39, pp. 170-210.

accepted, interpreted, and understood, but in few parts of the world will he be believed.

The boundaries between successive levels are not sharp and distinct. One shades off gradually into the next. Still the hierarchy is real enough and important to keep in mind. Simpler types of psycholinguistic processes can be studied rather intensively; already we know much about hearing and matching. Accepting and interpreting are just now coming into scientific focus. Understanding is still over the horizon, and pragmatic questions involving belief systems are presently so vague as to be hardly worth asking. But the whole range of processes must be included in any adequate definition of psycholinguistics.

I PHRASED the description of these various psycholinguistic processes in terms of a listener; the question inevitably arises as to whether a different hierarchy is required to describe the speaker. One problem a psycholinguist faces is to decide whether speaking and listening are two separate abilities, co-ordinate but distinct, or whether they are merely different manifestations of a single linguistic faculty.

The mouth and ear are different organs; at the simplest levels we must distinguish hearing and matching from vocalising and speaking. At more complex levels it is less easy to decide whether the two abilities are distinct. At some point they must converge, if only to explain why it is so difficult to speak and listen simultaneously. The question is where.

It is easy to demonstrate how important to a speaker is the sound of his own voice. If his speech is delayed a fifth of a second, amplified, and fed back into his own ears, the voice-ear asynchrony can be devastating to the motor skills of articulate speech. It is more difficult, however, to demonstrate that the same linguistic competence required for speaking is also involved in processing the speech of others.

Recently Morris Halle and Kenneth Stevens of the Massachusetts Institute of Technology revived a suggestion made by Wilhelm von Humboldt over a century ago.⁷ Suppose we accept the notion that a listener recognises what he hears by comparing it with some internal representation. To the extent that a match can be obtained, the input is accepted and interpreted. One trouble with this hypothesis, how-

ever, is that a listener must be ready to recognise any one of an enormous number of different sentences. It is inconceivable that a separate internal representation for each of them could be stored in his memory in advance. Halle and Stevens suggest that these internal representations must be generated as they are needed by following the same generative rules that are normally used in producing speech. In this way the rules of the language are incorporated into the theory only once, in a generative form; they need not be learned once by the ear and again by the tongue. This is a theory of a language-user, not of a speaker or a listener alone.

The listener begins with a guess about the input. On that basis he generates an internal matching signal. The first attempt will probably be in error; if so, the mismatch is reported and used as a basis for a next guess, which should be closer. This cycle repeats (unconsciously, almost certainly) until a satisfactory (not necessarily a correct) match is obtained, at which point the next segment of speech is scanned and matched, etc. The output is not a transformed version of the input; it is the programme that was followed to generate the matching representation.

The perceptual categories available to such a system are defined by the generative rules at its disposal. It is also reasonably obvious that its efficiency is critically dependent on the quality of the initial guess. If this guess is close, an iterative process can converge rapidly; if not, the listener will be unable to keep pace with the rapid flow of conversational speech.

A listener's first guess probably derives in part from syntactic markers in the form of intonation, inflection, suffixes, etc., and in part from his general knowledge of the semantic and situational context. Syntactic cues indicate how the input is to be grouped and which words function together; semantic and contextual contributions are more difficult to characterise, but must somehow enable him to limit the range of possible words that he can expect to hear.

How HE IS ABLE TO DO THIS is an utter mystery, but the fact that he can do it is easily demonstrated.

The English psychologist David Bruce recorded a set of ordinary sentences and played them in the presence of noise so intense that the voice was just audible, but not intelligible.⁸ He told his listeners that these were sentences on some general topic—sports, say—and asked them to repeat what they heard. He then told them they would hear more sentences on a different topic, which they were also to repeat. This was done several times. Each time the listeners

⁷ M. Halle and K. N. Stevens, "Speech Recognition: A Model and a Program for Research," *IRE Transactions on Information Theory* (1962), vol. 11-8, pp. 155-159.

⁸ "Effects of Context upon the Intelligibility of Heard Speech," in *Information Theory*, edited by Colin Cherry (Butterworths, London, 1956, pp. 245-252).

repeated sentences appropriate to the topic announced in advance. When at the end of the experiment Bruce told them they had heard the same recording every time—all he had changed was the topic they were given—most listeners were unable to believe it.

With an advance hypothesis about what the message will be we can tune our perceptual system to favour certain interpretations and reject others. This fact is no proof of a generative process in speech perception, but it does emphasise the important role of context. For most theories of speech perception the facilitation provided by context is merely a fortunate though rather complicated fact. For a generative theory it is essential.

Note that generative theories do not assume that a listener must be able to articulate the sounds he recognises, but merely that he be able to generate some internal representation to match the input. In this respect a generative theory differs from a motor theory (such as that of Sir Richard Paget) which assumes that we can identify only those utterances we are capable of producing ourselves. There is some rather compelling evidence against a motor theory. The American psychologist Eric Lenneberg has described the case of an eight-year-old boy with congenital anarthria; despite his complete inability to speak, the boy acquired an excellent ability to understand language.⁹ Moreover, it is a common observation that utterances can be understood by young children before they are able to produce them. A motor theory of speech-perception draws too close a parallel between our two capacities as users of language. Even so, the two are more closely integrated than most people realise.

I HAVE ALREADY offered the opinion that productivity sets the central problem for the psycholinguist and have even referred to it indirectly by arguing that we can produce too many different sentences to store them all in memory. The issue can be postponed no longer.

To make the problem plain, consider an example on the level of individual words. For several days I carried in my pocket a small white card on which was typed UNDERSTANDER. On suitable occasions I would hand it to someone. "How do you pronounce this?" I asked.

⁹ "Understanding Language without Ability to Speak: A Case Report," *Journal of Abnormal and Social Psychology* (1962), vol. 65, pp. 419-425.

¹⁰ A dog conditioned to salivate at the sound of a tone will also salivate, though less copiously, at the sound of similar tones, the magnitude declining as the new tones become less similar to the original. This phenomenon is called "stimulus generalisation."

He pronounced it.

"Is it an English word?"

He hesitated. "I haven't seen it used very much. I'm not sure."

"Do you know what it means?"

"I suppose it means 'one who understands.'"

I thanked him and changed the subject.

Of course, understander is an English word, but to find it you must look in a large dictionary where you will probably read that it is "now rare." Rare enough, I think, for none of my respondents to have seen it before. Nevertheless, they all answered in the same way. Nobody seemed surprised. Nobody wondered how he could understand and pronounce a word without knowing whether it was a word. Everybody put the main stress on the third syllable and constructed a meaning from the verb "to understand" and the agentive suffix "*er*." Familiar morphological rules of English were applied as a matter of course, even though the combination was completely novel.

Probably no one but a psycholinguist captured by the ingenuous behaviouristic theory that words are vocal responses conditioned to occur in the presence of appropriate stimuli would find anything exceptional in this. Since none of my friends had seen the word before, and so could not have been "conditioned" to give the responses they did, how would this theory account for their "verbal behaviour"? Advocates of a conditioning theory of meaning—and there are several distinguished scientists among them—would probably explain linguistic productivity in terms of "conditioned generalisations."¹⁰ They could argue that my respondents had been conditioned to the word understand and to the suffix—*er*; responses to their union could conceivably be counted as instances of stimulus generalisation. In this way, novel responses could occur without special training.

Although a surprising amount of psychological ingenuity has been invested in this kind of argument, it is difficult to estimate its value. No one has carried the theory through for all the related combinations that must be explained simultaneously. One can speculate, however, that there would have to be many different kinds of generalisation, each with a carefully defined range of applicability. For example, it would be necessary to explain why "understander" is acceptable, whereas "erunderstand" is not. Worked out in detail, such a theory would become a sort of Pavlovian paraphrase of a linguistic description. Of course, if one believes there is some essential difference between behaviour governed by conditioned habits and behaviour governed by rules, the paraphrase could never be more than a vast intellectual pun.

ORIGINAL COMBINATIONS of elements are the life blood of language. It is our ability to produce and comprehend such novelties that makes language so ubiquitously useful. As psychologists have become more seriously interested in the cognitive processes that language entails, they have been forced to recognise that the fundamental puzzle is not our ability to associate vocal noises with perceptual objects, but rather our combinatorial productivity—our ability to understand an unlimited diversity of utterances never heard before and to produce an equal variety of utterances similarly intelligible to other members of our speech community. Faced with this problem, concepts borrowed from conditioning theory seem not so much invalid as totally inadequate.

Some idea of the relative magnitudes of what we might call the productive as opposed to the reproductive components of any psycholinguistic theory is provided by statistical studies of language. A few numbers can reinforce the point. If you interrupt a speaker at some randomly chosen instant, there will be, on the average, about ten words that form grammatical and meaningful continuations. Often only one word is admissible and sometimes there are thousands, but on the average it works out to about ten. (If you think this estimate too low, I will not object; larger estimates strengthen the argument.) A simple English sentence can easily run to a length of twenty words, so elementary arithmetic tells us that there must be at least 10^{20} such sentences that a person who knows English must know how to deal with. Compare this productive potential with the 10^4 or 10^5 individual words we know—the reproductive component of our theory—and the discrepancy is dramatically illustrated. Putting it differently, it would take 100,000,000,000 centuries (one thousand times the estimated age of the earth) to utter all the admissible twenty-word sentences of English. Thus, the probability that you might have heard any particular twenty-word sentence before is negligible. Unless it is a cliché, every sentence must come to you as a novel combination of morphemes. Yet you can interpret it at once if you know the English language.

With these facts in mind it is impossible to argue that we learn to understand sentences from teachers who have pronounced each one and explained what it meant. What we have learned are not particular strings of words, but *rules* for generating admissible strings of words.

CONSIDER what it means to follow a rule; this consideration shifts the discussion of psycholinguistics into very difficult territory. The

nature of rules has been a central concern of modern philosophy and perhaps no analysis has been more influential than Ludwig Wittgenstein's. Wittgenstein remarked that the most characteristic thing we can say about "rule-governed behaviour" is that the person who knows the rules knows whether he is proceeding correctly or incorrectly. Although he may not be able to formulate the rules explicitly, he knows what it is to make a mistake. If this remark is accepted, we must ask ourselves whether an animal that has been conditioned is privy to any such knowledge about the correctness of what he is doing. Perhaps such a degree of insight could be achieved by the great apes, but surely not by all the various species that can acquire conditioned reflexes. On this basis alone it would seem necessary to preserve a distinction between conditioning and learning rules.

As psychologists have learned to appreciate the complexities of language, the prospect of reducing it to the laws of behaviour so carefully studied in lower animals has grown increasingly remote. We have been forced more and more into a position that non-psychologists probably take for granted, namely, that language is rule-governed behaviour characterised by enormous flexibility and freedom of choice.

Obvious as this conclusion may seem, it has important implications for any scientific theory of language. If rules involve the concepts of right and wrong, they introduce a normative aspect that has always been avoided in the natural sciences. One hears repeatedly that the scientist's ability to suppress normative judgments about his subject-matter enables him to see the world objectively, as it really is. To admit that language follows rules seems to put it outside the range of phenomena accessible to scientific investigation.

At this point a psycholinguist who wishes to preserve his standing as a natural scientist faces an old but always difficult decision. Should he withdraw and leave the study of language to others? Or should he give up all pretence of being a "natural scientist," searching for causal explanations, and embrace a more phenomenological approach? Or should he push blindly ahead with his empirical methods, hoping to find a causal basis for normative practices, but running the risk that all his efforts will be wasted because rule-governed behaviour in principle lies beyond the scope of natural science?

To withdraw means to abandon hope of understanding scientifically all those human mental processes that involve language in any important degree. To persevere means to face the enormously difficult, if not actually impossible task of finding a place for normative rules in a descriptive science.

Difficult, yes. Still one wonders whether these alternatives are really as mutually exclusive as they seem.

THE FIRST THING we notice when we survey the languages of the world is how few we can understand and how diverse they all seem. Not until one looks for some time does an even more significant observation emerge concerning the pervasive similarities in the midst of all this diversity.

Every human group that anthropologists have studied has spoken a language. The language always has a lexicon and a grammar. The lexicon is not a haphazard collection of vocalisations, but is highly organised; it always has pronouns, means for dealing with time, space, and number, words to represent true and false, the basic concepts necessary for propositional logic. The grammar has distinguishable levels of structure, some phonological, some syntactic. The phonology always contains both vowels and consonants, and the phonemes can always be described in terms of distinctive features drawn from a limited set of possibilities. The syntax always specifies rules for grouping elements sequentially into phrases and sentences, rules governing normal intonation, rules for transforming some types of sentences into other types.

The nature and importance of these common properties, called "linguistic universals," are only beginning to emerge as our knowledge of the world's languages grows more systematic.¹¹ These universals appear even in languages that developed with a minimum of interaction. One is forced to assume, therefore, either that (a) no other kind of linguistic practices are conceivable, or that (b) something in the biological makeup of human beings favours languages having these similarities. Only a moment's re-

fection is needed to reject (a). When one considers the variety of artificial languages developed in mathematics, in the communication sciences, in the use of computers, in symbolic logic, and elsewhere, it soon becomes apparent that the universal features of natural languages are not the only ones possible. Natural languages are, in fact, rather special and often seem unnecessarily complicated.

A popular belief regards human language as a more or less free creation of the human intellect, as if its elements were chosen arbitrarily and could be combined into meaningful utterances by any rules that strike our collective fancy. The assumption is implicit, for example, in Wittgenstein's well-known conception of "the language game." This metaphor, which casts valuable light on many aspects of language, can, if followed blindly, lead one to think that all linguistic rules are just as arbitrary as, say, the rules of chess or football. As Lenneberg has pointed out, however, it makes a great deal of sense to inquire into the biological basis for language, but very little to ask about the biological foundations of card games.¹²

Man is the only animal to have a combinatorially productive language. In the jargon of biology, language is "a species-specific form of behaviour." Other animals have signalling systems of various kinds and for various purposes—but only man has evolved this particular and highly improbable form of communication. Those who think of language as a free and spontaneous intellectual invention are also likely to believe that any animal with a brain sufficiently large to support a high level of intelligence can acquire a language. This assumption is demonstrably false. The human brain is not just an ape brain enlarged; its extra size is less important than its different structure. Moreover, Lenneberg has pointed out that nanocephalic dwarfs, with brains half the normal size but grown on the human blueprint, can use language reasonably well, and even mongoloids, not intelligent enough to perform the simplest functions for themselves, can acquire the rudiments.¹³ Talking and understanding language do not depend on being intelligent or having a large brain. They depend on "being human."

Serious attempts have been made to teach animals to speak. If words were conditioned responses, animals as intelligent as chimpanzees or porpoises should be able to learn them. These attempts have uniformly failed in the past and, if the argument here is correct, they will always fail in the future—for just the same reason that attempts to teach fish to walk or dogs to fly would fail. Such efforts misconstrue the basis for our linguistic competence: they fly in the face of biological facts.¹⁴

¹¹ *Universals of Language*, edited by J. Greenberg (M.I.T. Technology Press, Cambridge, Mass., 1963).

¹² E. Lenneberg, "Language, Evolution, and Purposeful Behavior," in *Culture in History: Essays in Honor of Paul Radin* (Columbia University Press, New York, 1960).

¹³ E. Lenneberg, I. A. Nichols, and E. R. Rosenberger, "Primitive Stages of Language Development in Mongolism," in the *Proceedings of the 42nd. Annual Meeting (1962) of the Association for Research in Nervous and Mental Diseases*.

¹⁴ The belief that animals have, or could have, languages is as old as man's interest in the evolution of his special talent, but the truth of the matter has long been known. Listen, for example, to Max Müller (*Three Lectures on the Science of Language*) in 1889: "It is easy enough to show that animals communicate, but this is a fact which has never been doubted. Dogs who growl and bark leave no doubt in the minds of other dogs or cats, or even of man, of what they mean, but

HUMAN LANGUAGE must be such that a child can acquire it. He acquires it, moreover, from parents who have no idea how to explain it to him. No careful schedule of rewards for correct or punishments for incorrect utterances is necessary. It is sufficient that the child be allowed to grow up naturally in an environment where language is used.

The child's achievement seems all the more remarkable when we recall the speed with which he accomplishes it and the limitations of his intelligence in other respects. It is difficult to avoid an impression that infants are little machines specially designed by nature to perform this particular learning task.

I believe this analogy with machines is worth pursuing. If we could imagine what a language-learning automaton would have to do, it would dramatise—and perhaps even clarify—what a child can do. The linguist and logician Noam Chomsky has argued that the description of such an automaton would comprise our hypothesis about the child's innate ability to learn languages or (to borrow a term from Ferdinand de Saussure) his innate *faculté de langage*.¹⁵

CONSIDER WHAT INFORMATION a language-learning automaton would be given to work with. Inputs to the machine would include a finite set of sentences, a finite set of non-sentences accompanied by some signal that they were incorrect, some way to indicate that one item is a repetition or elaboration or transformation of another, and some access to a universe of perceptual objects

growling and barking are not language, nor do they even contain the elements of language."

Unfortunately, Müller's authority, great as it was, did not suffice, and in 1890 we hear Samuel Butler ("Thought and Language," in his *Collected Essays*) reply that although "growling and barking cannot be called very highly specialised language," still there is "a sayer, a sayee, and a covenanted symbol designedly applied. Our own speech is vertebrated and articulated by means of nouns, verbs, and the rules of grammar. A dog's speech is invertebrate, but I do not see how it is possible to deny that it possesses all the essential elements of language."

Müller and Butler did not argue about the facts of animal behaviour which Darwin had described. Their disagreement arose more directly from differences of opinion about the correct definition of the term "language." To-day our definitions of human language are more precise, so we can say with correspondingly more precision why Butler was wrong.

¹⁵ N. Chomsky, "Explanatory Models in Linguistics," in *Logic, Methodology, and Philosophy of Science*, edited by E. Wägel, P. Suppes, and A. Tarski (Stanford Univ. Press, Stanford, 1962, pp. 528-550).

and events associated with the sentences. Inside the machine there would be a computer so programmed as to extract from these inputs the nature of the language, *i.e.*, the particular syntactic rules by which sentences are generated, and the rules that associate with each syntactic structure a particular phonetic representation and semantic interpretation. The important question, of course, is what programme of instructions would have to be given to the computer.

We could instruct the computer to discover any imaginable set of rules that might, in some formal sense of the term, constitute a grammar. This approach—the natural one if we believe that human languages can be infinitely diverse and various—is doomed from the start. The computer would have to evaluate an infinitude of possible grammars; with only a finite corpus of evidence it would be impossible, even if sufficient time were available for computation, to arrive at any unique solution.

A language-learning automaton could not possibly discover a suitable grammar unless some strong *a priori* assumptions were built into it from the start. These assumptions would limit the alternatives that the automaton considered—limit them presumably to the range defined by linguistic universals. The automaton would test various grammars of the appropriate form to see if they would generate all of the sentences and none of the non-sentences. Certain aspects would be tested before others; those found acceptable would be preserved for further evaluation. If we wished the automaton to replicate a child's performance, the order in which these aspects would be evaluated could only be decided after careful analysis of the successive stages of language acquisition in human children.

The actual construction of such an automaton is, of course, far beyond our reach at the present time. That is not the point. The lesson to learn from such speculations is that the whole project would be impossible unless the automaton—and so, presumably, a child—knew in advance to look for particular kinds of regularities and correspondences, to discover rules of a rather special kind uniquely characteristic of human language in general.

The features that human infants are prepared to notice sharply limit the structure of any human language. Even if one imagines creating by decree a Newspeak in which this generalisation were false, within one generation it would have become true again.

Psycholinguistics does not deal with social practices determined arbitrarily either by caprice or intelligent design, but with practices that grow organically out of the biological nature

of man and the linguistic capacities of human infants. To that extent, at least, it is possible to define an area of empirical fact well within the reach of our scientific methods.

ANOTHER LINE of scientific investigation is opened up by the observation that we do not always follow our own rules. If this were not so, of course, we would not speak of rules, but of the laws of language. The fact that we make mistakes, and that we can know we made mistakes, is central to the psycholinguistic problem. Before we can see the empirical issue this entails, however, we should first draw a crucial distinction between theories of language and theories of the users of language.

There is nothing in the linguistic description of a language to indicate what mistakes will occur. Mistakes result from the psychological limitations of people who use the language, not from the language itself. It would be meaningless to state rules for making mistakes.

A formal characterisation of a natural language in terms of a set of elements and rules for combining those elements must inevitably generate an infinitude of possible sentences that will never occur in actual use. Most of these sentences are too complicated for us. There is nothing mysterious about this. It is very similar to the situation in arithmetic where a student may understand perfectly the rules for multiplication, yet find that some multiplication problems are too difficult for him to do "in his head," *i.e.*, without extending his memory capacity by the use of pencil and paper.

There is no longest grammatical sentence. There is no limit to the number of different grammatical sentences. Moreover, since the number of elements and rules is finite, there must be some rules and elements that can recur any number of times in a grammatical sentence. Chomsky has even managed to pinpoint a kind of recursive operation in language that, in principle, lies beyond the power of any finite device to perform indefinitely often. Compare these sentences:

- (R) Remarkable is the rapidity of the motion of the wing of the hummingbird.
- (L) The hummingbird's wing's motion's rapidity is remarkable.
- (E) The rapidity that the motion that the wing that the hummingbird has has has is remarkable.

When you parse these sentences you find that the phrase structure of (R) dangles off to the right; each prepositional phrase hangs to the

noun in the prepositional phrase preceding it. In (R), therefore, we see a type of recurring construction that has been called right-branching. Sentence (L), on the other hand, is left-branching; each possessive modifies the possessive immediately following. Finally, (E) is an onion; it grows by embedding sentences within sentences. Inside "The rapidity is remarkable" we first insert "the motion is rapid" by a syntactic transformation that permits us to construct relative clauses, and so we obtain "The rapidity that the motion has is remarkable." Then we repeat the transformation, this time inserting "the wing has motion" to obtain "The rapidity that the motion that the wing has has is remarkable." Repeating the transformation once more gives (E).

It is intuitively obvious that, of these three types of recursive operations, self-embedding (E) is psychologically the most difficult. Although they seem grammatical by any reasonable standard of grammar, such sentences never occur in ordinary usage because they exceed our cognitive capacities. Chomsky's achievement was to prove rigorously that any language that does *not* restrict this kind of recursive embedding contains sentences that cannot be spoken or understood by devices, human or mechanical, with finite memories. Any device that uses these rules must remember each left portion until it can be related to its corresponding right portion; if the memory of the user is limited, but the number of admissible left portions is not, it is inevitable that some admissible sentences will exceed the capacity of the user to process them correctly.¹⁶

It is necessary, therefore, to distinguish between a description of the language in terms of the rules that a person *knows* and uses and a description of that person's *performance* as a user of the rules. The distinction is sometimes criticised as "psycholatry" by strict adherents of behaviourism; "knowing" is considered too mentalistic and subjective, therefore unscientific. The objection cannot be taken seriously. Our conception of the rules that a language-user knows is indeed a hypothetical construct, not something observed directly in his behaviour. But if such hypotheses were to be forbidden, science in general would become an empty pursuit.

Given a reasonable hypothesis about the rules that a language-user knows, the exploration of his limitations in following those rules is proper work for an experimental psychologist. "Psychology should assist us," a great linguist once said, "in understanding what is going on in the mind of speakers, and more particularly how they are led to deviate from previously existing rules in consequence of conflicting

¹⁶N. Chomsky, *Syntactic Structures* (Mouton, The Hague, 1957).

tendencies." Otto Jespersen made this request of psychology in 1924; now at last the work is beginning.¹⁷

ONE EXAMPLE. Stephen Isard and I asked Harvard undergraduates to memorise several sentences that differed in degree of self-embedding. For instance, the twenty-two words in the right-branching sentence, "We cheered the football squad that played the team that brought the mascot that chased the girls that were in the park," can be re-arranged to give one, two, three, or four self-embeddings; with four it becomes, "The girls (that the mascot (that the team (that the football squad (that we cheered) played) brought) chased) were in the park." One self-embedding caused no difficulty; it was almost as easy to memorise as the sentence with none. Three or four embeddings were most difficult. When the sentence had two self-embeddings—"The team (that the football squad (that we cheered) played) brought the mascot that chased the girls that were in the park"—some subjects found it as easy to memorise as sentences with zero or one embedding, others found it as difficult as sentences with three or four. That is to say, everybody can manage one embedding, some people can manage two, but everybody has trouble with three or more.

Records of eye movements while people are reading such sentences show that the trouble begins with the long string of verbs, "cheered played brought," at which point all grasp of the sentence structure crumbles and they are left with a random list of verbs. This is just what would be expected from a computer executing a programme that did not make provision for a sub-routine to refer to itself, *i.e.*, that was not recursive. If our ability to handle

this type of self-embedded recursion is really as limited as the experiment indicates, it places a strong limitation on the kinds of theories we can propose to explain our human capacities for processing information.

ON THE SIMPLER levels of our psycholinguistic hierarchy the pessimists are wrong; much remains there to be explored and systematised by scientific methods. How far these methods can carry us remains an open question. Although syntax seems well within our grasp and techniques for studying semantic systems are now beginning to emerge, understanding and belief raise problems well beyond the scope of linguistics. Perhaps it is there that scientific progress will be forced to halt.

No psychological process is more important or difficult to understand than understanding, and nowhere has scientific psychology proved more disappointing to those who have turned to it for help. The complaint is as old as scientific psychology itself. It was probably first seen clearly by Wilhelm Dilthey, who called for a new kind of psychology—a kind to which Karl Jaspers later gave the name "*verstehende Psychologie*"—and in one form or another the division has plagued psychologists ever since. Obviously a tremendous gulf separates the interpretation of a sentence from the understanding of a personality, a society, a historical epoch. But the gap is narrowing. Indeed, one can even pretend to see certain similarities between the generative theory of speech perception discussed above and the reconstructive intellectual processes that have been labelled *verstehende*. The analogy may some day prove helpful, but how optimistic one dares feel at the present time is not easily decided.

Meanwhile, the psycholinguists will undoubtedly continue to advance as far as they can. It should prove interesting to see how far that will be.

¹⁷ *The Philosophy of Grammar* (Allen and Unwin, London, 1924, p. 344).

Robert Lowell

Brunetto Latini

Canto XV of Dante's Inferno
(For Lillian Hellman)

And now we walked along the solid mire
above a brook whose fuming mist protected
water and banks from the surrounding fire.
Just as the men of Flanders threw up huge
earthworks to stop the sea that always threatens
their fields and cattle between Ghent and Bruges,
or Paduans along the Brenta spread
out dykes to shield their towns and towers against
spring thawing the Carinzian watershed—
on such a plan the evil engineer,
whoever he was, had laid his maze of dykes,
though on a smaller scale, and with less care.
By now we'd gone much deeper underground,
and left the bleeding wood so far behind
I'd have seen nothing, if I'd looked around.
We met a company of spirits here,
trooping below us on the sand. Each one
stared closely at our faces—as men peer
at one another under the new moon,
or an old tailor squints into his needle,
these puckered up their brows and glowered. Soon,
I saw a man whose eyes devoured me, saying,
“This is a miracle.” He seized my sleeve,
and as I felt his touch, I fixed my eyes
with such intensity on his crusted face,
that its disfigurement could not prevent
my recognising who he was. “Oh, Oh,”
I answered groaning, as I stretched my hand
to touch his arm, “are you here Ser Brunetto?”
He answered, “Do not be displeased, my Son,
if Brunetto Latini turn and walk a little