Invited address, Division 3, APA 1970

Perceptual Learning and the Theory of Word Perception

Eleanor J. Gibson

Cornell University

When the Program Committee of Division 3 honored me with the invitation to address this group today, I was told "feel free to talk on any aspect of your current research or theory." I pondered the choice for some time and concluded that I wanted to talk about both research and theory. I have spent the last 10 years developing a theory of perceptual learning; but along with that, I have followed a program of research on reading. This research was not generated by specific classroom problems, but by a desire to understand the learning process involved. It is theory-based, and thus it is related to my notions of perceptual learning. I have decided to try, today, to show how the two are related. More specifically, since the whole topic of reading is too big, I am going to talk about perceptual learning and the theory of word perception. Most of my recent experiments originated in this setting.

I shall begin by summarizing my theory of perceptual learning. I intend to do this as briefly as possible, because some of you (I hope) have read my book about it. I divide the theory into attempts to answer three questions: First, what is learned? Second, How? What are the processes involved? Third, what is the motivation and reinforcement for perceptual learning?

1

What is Learned?

I believe that what is learned in perceptual learning are distinctive features of things, of representations of things, and of symbolic entities like words; also the invariants of events that occur over time; and that we learn economical structuring of both. I think the information for learning these is potentially present in stimulation, to be picked up by the observer given the proper conditions for it.

invariants f. ex. a recurring phoneme in the flow of speach

structures = relations between features to be discovered because they exist in the world

Let me give some examples. Sets of distinctive features characterize objects and entities both natural and artifactual--the furnishings of the world, such as people, dwelling places, things to eat; and, particularly relevant for my present topic, symbolic entities written on pieces of paper, like letters and words. The set of letters of our alphabet is characterized by a set of distinctive features, which in different combinations permit a unique characterization of each one. My students and I have spent much time trying to describe the set of distinctive features that are shared by letters of the Roman alphabet. We have had some success, since confusion matrices obtained experimentally reveal, via cluster analysis, some contrasting features that can be diagrammed in a tree structure.

[The slide shows you the feature-splits obtained from an analysis of discrimination data of 7-year-old children for a set of 9 letters.]

What about invariants of events? These occur over time. The nicest examples of learning to detect them occur in perceptual development. The constancies can only be understood as invariants over transformations occurring in an event like approaching or receding of an object, or rotation of it. Perceived existence of an object despite temporary occlusion by a screen of some sort (the event of "going in front lf" or "going behind") is another. What has this to do with words? Events like appearance, disappearance, and reappearance have meaning, and these very meanings are expressed spontaneously in the child's earliest two-word utterances (Hi Daddy; all-gone ball; more milk; ball again). These utterances appear to be invoked by the event, quite literally mapped to it (Brown, 1970). Detection of invariance, in other words, is prior to symbolic referential meanings and is reflected in them. Semantic features of words generally indicate perceivable features of the world, both things and events. Verbs for instance can be classified as action or state, an important semantic distinction arising directly from differentiating invariants in the environment. (?) Nouns can be classified as count or mass, a distinction that depends on differentiating things with fixed contours from those that are fluid, like sand or water.

Finally, what are some higher-order structures discovered by perceptual learning? Entities in the world, both natural and artifactual, have structure; that is, relations between features. These relations can be subordinate and superordinate. Both superordinate and subordinate structure are progressively discovered in development, I think (for both objects and events--the structure of events was referred by by Brunswik as the "causal texture of the environment"). Examples of the subordinate and superordinate structure of words are obvious. Words come in sentences; sentences are parts of paragraphs; phrases are parts of sentences.

Processes of Perceptual Learning

How are these things learned? Not, I believe by associating a response or an image or a word to a "stimulus." Distinctive features and invariants must be discovered--extracted from the multiplicity of information in stimulation that contains adjacent and temporal order. We have always accepted the notion of abstraction to explain the genesis of concepts. I think a process akin to abstraction--the dissociation of an 1.
invariant from transforming or variable context--happens also at a perceptual level. The phoneme is segregated from the flow of speech heard by the infant and its invariant distinctive features are abstracted from many varying samples and over many transformations produced by different speakers. The process has to be one of abstraction--there is no response to be associated, nor is there an identically repeated unvarying stimulus.

What happens to the variable, irrelevant components of stimulation? When something invariant is abstracted and selected out for attention, what happens to the rest? I think that the process of abstraction is accompanied by a filtering process that attenuates and suppresses the 2.
irrelevant--this is what happens to the words that aren't heard in the dichotic listening experiments, for instance.

Finally, a very important process in perceptual learning is the operation of mechanisms of external attention. The sensory systems are all active and exploratory. "Looking," "listening," and "feeling" are terms that describe the search for information in stimulation. They also underline the fact that perception and perceptual learning are active processes. There is improvement and flexibility in attentive strategies depending on age and the business in hand, as I hope to show.

Motivation and Reinforcement

What selects the good strategy, the economical feature, the structure that most effectively orders the information? This is the question of motivation and reinforcement in perceptual learning. Again, I am looking to new concepts. I do not think the motivation is to be found in drives like tissue needs, nor the reinforcement in reduction of a metabolic drive nor in cessation of punishment. I think there is a built-in need to get information about one's environment. One could call this, as a number of psychologists do, "intrinsic cognitive motivation." I can also (and do) call it the "search for invariance." This motivation is always re- If there is such a search
lated to the task in hand, for different information is needed for different tasks. But active looking for information about the solid, permanent safe places of the world, and the invariant aspects of events (like the swift approach of an object) is essential for behaviors like locomotion and avoidance of predators. A need to understand what others are saying seems equally basic for learning to comprehend language.

Reinforcement of perceptual learning (indeed of all cognitive learning) is, I would contend, the reduction of uncertainty. I believe that discovery of structure, or discovery of the economical distinctive feature, or of the rule describing an invariant reduces the "information load," leads to permanent perceptual reorganization for the viewer of the world so viewed, and leads to repetition of the successful strategy when a similar occasion occurs. My evidence for this is slight, but I have been trying to conduct experiments to see whether discovery of structure is indeed reinforcing.

Let me describe one that is in progress. Children are presented with the task of solving anagrams. They are simple ones, all rearrangements of letters in words encountered by the end of third grade. Five of them are presented on a board, one above the other, each covered with a strip of ribbon which the child is allowed to remove, one anagram at a time. The letters are engraved on bits of wood which are magnetized so that they stick in place. The child is allowed to rearrange them as he tries to solve the anagram.

Now, for one group of Ss, the five anagrams are arranged in an order that makes a sentence, like the following:

Jill
helps
Mom
clean
house

The child solves each anagram in turn. When all five are done, that board is removed and he is given a new one. Again, the five anagrams, when solved, make a sentence. He solves six such sets of anagrams. The hope is that he will discover that the anagrams, when solved, make a sentence. If he does, will he try to use the potential structure on later presentations?

This group is compared to two other groups. One of them has the same anagrams, but they are arranged on the board in scrambled order so there is no sentence. The other starts with the anagrams in sentence order, but after the first four boards have been done, the last two are presented in the scrambled order. If these children discovered the structure, and if it had a reinforcing effect, they should now show interference in solution compared to the control group. The group with all six trials presented in sentence order should show a rising curve for speed of solution, compared to the control group with scrambled order.

To summarize, perceptual learning is active and is motivated by a search for information, especially invariant features. Learning progresses toward the optimal strategy, picking up the most economical structure and features.

How are Words Perceived?

Now I shall tackle the question of how words are perceived. (I am particularly interested in how they are perceived in reading, since that is where my research has been centered.) My answer sounds simple: Words, like other entities in our environment, are perceived by detecting their distinctive features. But what are a word's distinctive features? A word is part of a vast system of information. To answer this question, we have to identify this information.

Features of Words

The best way to represent it is to refer to what is learned when we learn to hear and speak, and to read and write. I have summarized this in a chart. Three aspects of the information dealt with in hearing-speaking (phonological, semantic and syntactic information) are presented on the left-hand side. Three parallel aspects for reading-writing (graphological, semantic, and syntactic information) are presented on the right. The relations between these are of great interest when one asks how we learn to read, but I shall concentrate now on the question of what information is picked up in reading a word, and then, tentatively, how it is picked up.

Examination of this chart tells me what classes of features a word can have. It can have graphological features of many kinds. It has a characteristic shape and length (referred to as "word shape" by reading teachers). Within the word there are letter shapes, themselves differentiated by distinctive features. And then the word has orthographic structure; letters are combined into words according to a rule system so that given combinations or clusters are permissible only in certain locations and contexts. Q must be followed by U, for instance; Qu can begin a word or a syllable, but it cannot end them.

How do we know whether these potential features of words are actually being detected as features? A useful method has been to set up tasks that produce some errors and then to study what is confused in the errors. We can infer some of the features of letter shapes that are being noticed when a child confuses E and F but not E and O. We infer that a larger graphic structure is perceived when the child confuses palindromes, like "saw" and "was." He is generalizing something relational. We can also study features by looking at accurate generalizations, of course, like discovering a rule about spelling patterns and transferring it to new cases (see Gibson, 1970).

A word also has potential phonological features, even when it is read. When we read poetry, we are keenly aware of acoustic similarity in the rhymes, for instance. Homophones, albeit they are spelled differently, are sometimes confused. The fact that a string of letters presented visually is pronounceable vastly affects its readability. Some years ago I demonstrated this with former students of mine, Anne Pick, Harry Osser and Marcia Hammond (1962). We made up pseudo-words that could

Esempi

?

?

Esempio

be pronounced, though they were not real words and were low in meaningfulness. We constructed them so that they began with initial consonant clusters that could not end a word, and were terminated with clusters which could not begin one.

E.g., GLURCK
CKURGL

Unpronounceable strings were constructed from the pronounceable ones by exchanging the initial and final consonant clusters.

When these letter-strings were presented visually, with a short exposure, subjects invariably read the pronounceable version more accurately than the unpronounceable version. The advantage might have been due to easier pick-up of acoustic structure; or easier pick-up of articulatory structure. But it could also have been influenced by good orthographic structure, quite apart from phonological features, for we found later that profoundly deaf Ss showed the same relative advantage. These are all potential features of a word, but from different feature classes.

A word has many semantic features. It has semantic markers of various kinds indicating classes or properties, like "proper" nouns, "count" nouns, "mass" nouns. It may stand for objects belonging to categories like edible things, or for events like looming or disappearance. It has similarity and contrast relations--that is, synonyms and antonyms. We know these features are picked up because one can show experimentally that they are confusible. And one can show pick-up of semantic features by "clustering" in recall. Semantic features of a word include values, as well. Words can be rated for pleasantness-unpleasantness, or ranked on Osgood's semantic-differential scales.

A word's <u>syntactic</u> features are equally obvious. It is a part of speech, like noun or verb or adjective. It has a role in a sentence, like subject or object. And it may possess a morphological marker, like pluralization or tense. This last feature, morphological information, was investigated in an experiment by Lynne Guinet and myself.

We wondered whether a <u>longer</u> word could be perceived tachistoscopically when length was increased by adding a well-known inflected ending to a base word. Are the endings themselves a unitary structure?

We added inflected verb endings to three types of base words; real familiar words, pronounceable pseudo-words that were anagrams of them, and unpronounceable pseudo-words. There were also four lengths of base-word. The slide shows all the word-types and the endings used. These words--base-words and inflected words--were shown tachistoscopically to subjects from third grade, from fifth grade, and from the elementary psychology course. The results were not just what I had anticipated, but they were very illuminating. Subjects did <u>not</u> perceive a longer word correctly when a base word was expanded by adding an inflected ending. But the endings gave evidence of unitariness, for there were significantly fewer errors in inflected endings than in endings of base words of equivalent length (especially when these were not meaningful or pronounceable). Furthermore, when there were errors in the inflected endings, the errors tended to be substitutions of other inflected endings. There was a confusion within a morphological feature-class. These latter two tendencies increased from third to fifth grade, showing progressive pick-up of syntactic features.

Why wasn't a longer word perceived when the syntactic marker was added? Because, as I should have realized, the subject had to process an extra feature. It is features that must be processed, not elements like letters or syllables, and the process takes time. No matter how long the base word, its features must be recognized and so must the morphological tag.

Three Hypotheses

I would like to suggest three hypotheses about how these features are perceived. First, I think they are processed independently and sequentially, in a layered fashion, a kind of hierarchy. If pick-up time is cut short, a feature at the bottom of the hierarchy may be missed. I don't think it is deleted, but just not fully detected, so a confusion within its feature-class may result (like substituting ing for ed in the experiment I just described).

Proof-reader's errors are such a case, but as a matter of fact, spelling is generally noticed in reading, when the time is not limited. Something looks odd and we go back after a sentence or two and verify the mistake.

Another familiar piece of evidence for ordered, independent pick-up of word features is "semantic satiation" or "loss of meaning." If I put a printed word in front of you and tell you to stare at it for five minutes, its meaning gradually slips away. First the semantic features go, then the phonological features go, and one is left finally with the graphic features and even these will eventually fragment. There is a very interesting point here. Meaning, for an adult reader, is embedded in the word. He doesn't begin by decoding it letter by letter; the concept symbolized by the word "hits him." It is specified in stimulus information.

The Stroop test is further proof of this. In this test, two features are put in opposition, meaning and a graphic feature, the color of the ink the word is printed in. The subject is shown an array of words and asked to name the color each word is printed in, going from left to right as in reading--green, red, blue and so on. But the words themselves are the names of colors. When the name and the color of the ink conflict, the subject is in trouble. The name comes first to mind and his performance is badly slowed up compared to just giving the name of a color patch. There seems to be less interference in this task with very young readers. The meaning isn't yet as firmly embedded in, or specified by, the word on the page.

The second hypothesis is that there is a developmental change with age and schooling in feature analysis and pick-up. We can find evidence for this in reading errors at progressive levels of instruction. Errors in oral reading through the first grade were studied by Rose Marie Weber (1970) and Andrew Biemiller (1968). The earliest errors, at the beginning of first grade, can generally be attributed to meaningful context. When the child reaches a word he does not "know," he uses all the semantic information at his disposal (context of the words already decoded, pictures, and so on) and guesses. He produces a word that makes sense, but bears no resemblance otherwise to the one on the page. A little later the child stops when he reaches an unknown word and simply says nothing. This is a transition to a stage where errors become determined by graphic similarity. The child is engrossed by discriminating letters and by correspondences of letters and sounds. Semantic features of the word are temporarily lowered

7 in priority as the child strives to "break the code." (This is the period

when he may be chided by the teacher for "reading without expression," but the stage nevertheless marks progress.)

Semantic features return as the decoding process becomes easier and the orthographic features demand a lesser share of the child's attention. The child reads with more "expression." But the various rule systems (orthographic and syntactic) probably do not operate fully as important structural constraints until later. The influence of orthographic structure begins to be quite apparent by third grade. Besides the evidence from my tachistoscopic experiments (Gibson, Osser, & Pick, 1963), it has been shown by Rosinski and Wheeler (personal communication) that third graders can judge correctly 90 percent of the time which of two pseudo-words is "more like a word." This judgment is not affected by length of the letter string. First graders, however, make this judgment at a chance level.

In the experiment I reported, morphological inflections such as verb endings were found to operate as unitary features of a written word. This function was more apparent in the fifth graders than the third graders, and most apparent in college students.

Finally, syntactical constraints like phrase structure and grammatical conventions--the word's role in the sentence--have been shown in studies with the eye-voice span to operate in reading, and again this rule system increases in usefulness as reading skill progresses, being noticeably more functional in fifth than in third graders. The influence of phrase boundaries, for example, only shows up after grade 2 (Levin and Turner, 1968). The young reader does appear, then, to show a developmental sequence in pick-up of word-features. This progression is no doubt not as fixed as I may have implied and probably begins before very long

?

?

to be influenced by the reader's task, in accordance with my next hypothesis.

My third hypothesis holds that the order of pick-up of word-features changes with the task. To put it a little differently, priorities of pick-up are geared strategically to task utility. I repeat now my earlier point that perception is an active search for information and that the perceptual strategy that develops will be as economical as possible. Common sense suggests that this is so--when we are looking for a weather report in the newspaper, we assign a very low priority to graphic features of the words we read. But in addition there exists a large number of experiments which go to prove my point, some of my own and many by others.

We can influence priorities by instructions, of course. In a tachistoscopic experiment, if a subject is told to guess at words shown him, features like meaning and word frequency are evident. If he is told to report only what is literally seen graphic features are advanced in priority. But quite aside from instructions, different tasks seem to have acquired their own priorities in the cognitive economy either in the course of development or by learning during the task. All a word's features have their importance for one task or another.

Phonological features, we have learned in recent years, have a high priority for short-term memory. If I am trying to hang onto a telephone number and someone speaks to me before I have dialed it, it is lost. Conrad (1964) and others have shown in a number of experiments that acoustic similarity produces confusions in short-term memory even when the material is presented visually (Baddeley, 1968). Graphic and semantic confusions in short term memory, on the other hand, are infrequent

(Baddeley, 1966). Articulatory similarity may play a role in short-term memory tasks, as well (Crowder & Morton, 1969). I shall class this with phonological information, but it undoubtedly plays its own role, distinct from acoustic similarity, in some tasks.

A word's pronounceability, I suggested earlier, is one of its phonological features. It is one which strongly facilitates pick-up of tachistoscopically-presented displays. Does it do so equally for another task, like later recall? I tried (Gibson, Bishop, Schiff, & Smith, 1964), to separate pronounceability and semantic reference and to compare their effects in these two kinds of tasks. Trigrams were prepared which either rated high in pronounceability (like MIB), or in referential meaning (like IBM), or in neither (like MBI). In one experiment, they were presented tachistoscopically and recognition thresholds were obtained. Pronounceability very effectively facilitated accurate perception of the trigram. Meaning helped a little, compared to the items low in both pronounceability and meaning, but much less so. In another experiment, the same items were presented to subjects for two seconds each and 24 hours later recall was tested. This time the effect of meaning and pronounceability was reversed. Meaning facilitated recall far more than pronounceability and there was evidence of categorizing, for some subjects made up sets of initials, like FDR, that had not been in the list.

?

While phonological features of words dominate pick-up in short-term memory, there is considerable evidence to show that they have low priority for long-term memory. Acoustic similarity has little or no interfering effect in a paired-associate retroactive-inhibition paradigm (Dale and Baddeley, 1969; Bruce and Murdock, 1968), whereas semantic similarity does. Wickens, Ory and Graf (1970) found that acoustic similarity had some negative effects in a transfer paradigm, but the effect was extremely slight compared to semantic similarity in the same paradigm. Sharing taxonomic category membership was a powerful influence on the subject's ability to recall items of a list for either good or ill depending on task relations.

The utility (and utilization) of semantic features of a word for later recall has often been demonstrated by evidence of clustering in long-term recall. ? This was brought out cleverly in a recent experiment by Hyde and Jenkins (1970). In this experiment, the subjects were sometimes asked to do two tasks at once. They were presented with a list of words (for later recall). In two conditions, Ss had to extract some graphic information about a word as it was presented--either estimate its length (number of letters), or detect the presence or absence of the letter E. Another group had to rate the word as it was presented for pleasantness or unpleasantness.

Compared to a control group with no second task, recall was greatly reduced for the first two groups and so was the amount of organization in recall as measured by clustering of words in categories. But the task of rating words as pleasant or unpleasant did not reduce recall nor organization in recall as compared with a control group that had no incidental task. When a subject is performing a second task that gives priority to semantic features of the word, neither recall nor its organization suffer. But when the second task requires attention to a word's graphic features, like detecting e's or estimating word-length, the semantic pick-up which is apparently vital to free recall of words is blocked.

The Hyde and Jenkins experiment suggested to me that features of the same type, like semantic features of all kinds, are picked up in parallel, while different feature-types are processed sequentially. The value of the word--pleasant or unpleasant--could apparently be assessed at the same time as pick-up of semantic categories of the kind that operate in clustering. I will consider this again in connection with a different task, visual search.

Over the past five years, I have been particularly interested in visual search tasks, partly because I am interested in how perceptual search develops and also because they offer a good opportunity for studying what the subject learns incidentally. A lot of perceptual learning goes on during this task. Visual search also provides a fine opportunity for comparing pick-up of different types of word-features. I have made such comparisons in a number of experiments. The task is similar to one used by Ulric Neisser and involves scanning systematically through a matrix of letters for a target letter or word.

When the subject is asked to search for a letter target embedded in a context of other letters by scanning down a column of letters arranged 5 or 6 to a row, he very quickly sets his priorities for graphic features. Even 7-year-old children do this. Albert Yonas and I compared the effect of high and low graphic similarity of context letters to the target letter. Graphic similarity slows the scanning rate enormously, both for adults and children. Given an opportunity for practice, the subject will learn to scan for a single very distinctive feature, as Albert Yonas and Franklin Schapiro have both shown in appropriate transfer experiments.

What about phonological features of the letters in this task? They seem to be virtually, if not entirely ignored. Yonas and I tried to produce interference by exposing the subject to a voice pronouncing letters that sounded like the target letter while he scanned the list visually. There was no effect at all on scanning rate, even in children of seven years. Kaplan, Yonas, and Shurcliff (1966) compared the effect of high and low acoustic similarity of context letters to the target. That is, the target was embedded in a context of letters that rhymed with it (B, V, D, and so on) or in a context of letters that sounded unlike it. Acoustic similarity did not slow scanning rate at all, in contrast to the powerful effect of graphic similarity.

Subjects typically remember context letters in this task very poorly, even when they are tested by recognition with the letters presented to them visually (Schapiro, 1970). I wondered whether introduction of some structure of a higher order in the context would not bring it to the fore perceptually. If subjects can learn to take advantage of the most economical possible single graphic feature that distinguishes the target, will they discover and use orthographic structure if it is present?

Yvette Tenney and I performed an experiment in which we compared scanning rate for a letter target, either embedded in letter-strings which, though not meaningful, were orthographically possible words and were pronounceable; or with the target embedded in strings of the same letters scrambled so as to be unpronounceable. I had thought that orthographic structure might be picked up in parallel with graphic structure, which is so insistent in this task. If so, it might facilitate search, because the subject could filter the irrelevant context in larger units--strip it off in bigger chunks, so to speak.

On the other hand, if the subject subvocally articulated the pronounceable context items, scanning rate ought to be slowed down. (I have been trying to find out for a long time whether orthographic structure and pronounceability are necessarily tied.) We ran 76 children in the fifth grade in this experiment. There was no difference in mean rate of scan between the two conditions. What was happening? It would seem that the children could not, then, have been articulating the pronounceable items. Like lower-order phonological features, pronounceability in its literal sense of pronunciation seems not to influence the kind of verbal processing that goes on in this task. But what about the orthographic structure as such? Does it go unnoticed or can it not be used without an accompanying act of articulation that would be uneconomical? I think the latter may be the case here. When the children were questioned after the experiment they appeared, when context strings were pronounceable, to have been aware of it. Sometimes they commented on it spontaneously. But still it did not speed up the mean scanning rate. The child could of course have processed the whole string as a unit, and then searched it for the target letter as a second step. This would contradict the suggestion I made earlier, that a given class of word features gets processed simultaneously. ? It is also contradicted by recent experiments of Reicher (1969) and Wheeler (1970).

On the other hand, the children may have found the orthographic structure early in the task (since they so often commented on it), but learned to disregard it, perhaps because they could not use it without articulation--a handicap in a scanning task. This would be a kind of perceptual learning involving an inhibitory or filtering process, one of the three processes for perceptual selectivity that I hypothesized earlier. It may be that adult readers can use the structure without articulation and would speed up their scanning rate even though the children did not.

What would happen if semantic structure were introduced in the scanning task? Would it be picked up, when presented incidentally? Would it speed the search, and if so would it transfer as a strategy--that is, lead to a search for similar structure in a second task? I investigated these questions in another experiment with Yvette Tenney. We introduced semantic structure by building categorical meaning into the context to be searched through in looking for a target word.

Context words, in one condition, were all the same length and all belonged to the same semantic category--e.g., kinds of fruit. We used categorical material that had previously been shown to cluster well in recall experiments. The target word was the name of an animal, and it varied from trial to trial. The S was told to search for the name of an animal, rather than for a specified word, because we found in exploratory work that S searched for nothing but graphic features if he was given a specific target word. We wanted to force pick-up of the categorical relation among context words, if it could be done. Exper.

A control group had context words chosen at random as regards meaning, but equated with the other condition for frequency, length, and as many graphic attributes as possible. There were two other conditions in which a second scanning task with different context words followed. There was no common category of context words in these second tasks, since we thought an interference effect might appear for the Ss who used categorical structure in the first task. They should be slowed down by searching for a non-existent category.

The results of this experiment ruled out unequivocally the possibility that categorical meaning plays any role at all in a search task of this sort. Semantic structure of the kind we introduced is evidently not an economical feature for a search task. It did not speed up scanning rate and was not used by the Ss. Many Ss did not even notice the categorical relation within the context although they were told to use category membership for locating the target.

What they actually did, it turned out, was very economical indeed. They looked, after a few trials, not for a word of any particular meaning, but simply for any combination of letters that had not appeared before. Performance on the second task was indistinguishable for the groups compared, no matter whether there had been semantically structured context before or not.

Subjects do learn in this task. They learn the strategy that is most economical for the task. The conclusion is inescapable that semantic features of words have little utility for a search task and are ignored in favor of graphic features that do. It is not that structure is never utilized; graphic structure in the sense of a redundant graphic feature that helps differentiate the target is picked up and used as Franklin Schapiro has recently (1970) shown. But perceptual strategy in this task sets the graphic features a high priority, and all other features--semantic, acoustic, even orthographic, low.

One may ask, finally, if there is any task where syntactic features of words have priority. In laboratory tasks as far as I know, this does not seem to be the case, although they are certainly picked up, as Lynne Guinet and I showed. A feature like part of speech--noun vs. verb or adjective--is not very effective for producing clustering in recall nor generalization in verbal conditioning. That may be, one could object, because syntactical information is generally spread over a string of words and is imperfect in one. But although a word always tells us more in context, it appears that its different features are picked up independently and differentially. I find that there is one real-life task in which I seem, willy-nilly, to give first priority to syntactic information. That is, reading students' papers, or a thesis. A split infinitive, or a singular verb following the word "data," distracts me so that I lose the meaning!

Conclusion

Does perceptual learning occur in word perception? I think it does--both during development and within a task. Words contain many kinds of information, and we learn to perceive them as a complex of features. Words should not be thought of as made up of elements of a given length, but rather as entities possessing information classifiable as phonological, graphic, semantic, and syntactic features. All these kinds of information are in them, I think, in the sense that a word <u>specifies</u> its information. The perceiver <u>does</u> something, indeed, but he does not invent the information.

Word perception, like other kinds of perception, is active, searching for the relevant information in stimulation. Perceptual learning with words, like other examples of perceptual learning, develops toward the strategy that is most economical. This means that priorities for features shift, with practice in a task, toward those that have most utility for it.