INTERACTIVE PROCESSES IN READING

2 Interactive Processing Through Spreading Activation

David E. Rumelhart and James L. McClelland
Program in Cognitive Science
University of California, San Diego

A central issue in the development of a model of the reading process is the way the context in which a linguistic element is found affects the way that element is processed and ultimately interpreted. Rumelhart (1977) proposed a very general model, dubbed an *interactive model*, of the reading process that was designed to explicate the role of context during reading. In brief, an interactive model is one in which data-driven, bottom-up processing combines with top-down, conceptually driven processing to cooperatively determine the most likely interpretation of the input. Roughly speaking, processing in an interactive model of reading proceeds in the following way: The reader begins with a set of expectations about what information is likely to be available through visual input. These expectations, or initial hypotheses, are based on our knowledge of the structure of letters, words, phrases, sentences, and larger pieces of discourse, including nonlinguistic aspects of the current contextual situation. As visual information from the page begins to become available, it strengthens those hypotheses that are consistent with the input and weakens those that are inconsistent. The stronger hypotheses, in turn, make even more specific predictions about the information available in the visual input. To the degree that these hypotheses are confirmed, they are further strengthened, and the processing is facilitated.

Interactive processing is thus a form of cooperative processing in which knowledge at all levels of abstraction can come into play in the process of reading and comprehension. In earlier work on this topic, Rumelhart (1977) proposed a processing mechanism of this general sort. The mechanism was very general and suffered from a lack of direct connection to an empirical base and a related lack of specificity of exactly how the brain might actually carry out such complex computations. In this chapter (and the more complete descriptions of McClelland

& Rumelhart, 1980; and Rumelhart & McClelland, 1980), we have attempted to propose a mechanism that is both closely tied to an empirical data base and that is sufficiently specific to compare the results of computer simulations of the model with the empirical data.

Clearly, it is premature to develop a model of the kind of specificity we are proposing for the entire reading process. We have attempted to make a step in that direction by developing a model in which knowledge about words plays a central role in the perception of sequences of letters, which may or may not form familiar words. There is a large amount of literature on this topic, so there are a large number of important findings to constrain our theorizing. We view this domain as a kind of testing ground for the general question of how contextual information and general knowledge guide processing, but also as a testing ground in which we can get firm empirical predictions. We begin by reviewing some of the findings that we take as central to the area of word perception, and in so doing, we begin to suggest the outlines of an interactive model that can account for them.

SOME IMPORTANT FACTS ABOUT WORD PERCEPTION

Study of the perception of letters in words goes back a very long time in psychology, to the early research of Cattell (1886) and others (see Huey [1908] for a review). Most of these early experiments demonstrated that words could be perceived under conditions that were not sufficient to allow accurate perception of all the letters. In the late 1940s and 1950s, researchers in the "new look" in perception focused attention on the fact that more frequent words could be perceived more accurately than less frequent words at a given exposure level and could be perceived at lower exposure levels (Howes & Solomon, 1951). In addition, the important discovery was made that letters were perceived more accurately in nonwords that conformed to some of the statistical properties of words than in nonwords that did not (Miller, Bruner, & Postman, 1954).

Although these phenomena are consistent with the view that perception of the letters in words (and nonwords) is not merely a matter of independent perception of unrelated elements, it was possible to interpret them in terms of postperceptual guessing or forgetting processes. A large number of unrelated letters might pose a memory load that would limit accuracy of identification, even if all the letters were actually perceived correctly. And, it was argued, knowledge about what makes a word in English could allow subjects to *guess* imperfectly perceived letters in words more accurately than letters in unrelated strings.

In 1969, Reicher introduced a procedure that controls for these "nonperceptual" interpretations of the classic findings in word perception. In this procedure, the subject views a target item, followed by a mask, and is then given a pair of alternative letters, either of which might have been presented. For example, if the

item *WORD* is shown, the subject might be tested with a choice between *D* and *K*. On another trial, a nonword such as *ORWD* might be shown, followed by a choice, as before, between *D* and *K*. Subjects were more accurate in their choices when the item formed a word, even though the context letters really do not provide any clue as to which of the two letters, *D* or *K*, should be correct. Thus, the fact that the letters formed a word seemed to permit the subject to perceive more information relevant to a discrimination between the letters in the word. The possibility that postperceptual forgetting processes were responsible for the effect was eliminated by the fact that Reicher obtained a word advantage not only over letters in four-letter nonwords but also over letters presented in isolation.

Reicher's experiment has stimulated a large body of research, and his method has become the method of choice for studying the role of word contexts in facilitating perception of letters. For this reason, results obtained using this paradigm have been the primary focus of our modeling effort, although we do have occasion to note a few findings, which have been obtained using the classical whole report technique, that are consistent with our model.

The literature provides several important clues to the processes at work in word perception that have been central to the development of our model. One central finding is the fact that it is not necessary for words to be presented in a familiar visual form to produce a perceptual advantage over sequences of unrelated letters. McClelland (1976) found that the advantage of words over unrelated letters could be obtained using mixed upper- and lower-case type. Adams (1979) reproduced this result using stimuli in which the upper- and lowercase letters differed widely in size, so that the visual configuration was quite unfamiliar indeed. Mixing upper- and lowercase type does have a disruptive effect, and it appears to apply equally well to both words and nonwords when both types of stimuli are tested at comparable performance levels (Adams, 1979). These results suggest that an adequate model of word perception cannot rely (at least not exclusively) on recognition of familiar shapes of words to account for superior perception compared to unrelated letter strings. It would seem that it is knowledge of the arrangement of letters in familiar words that facilitates their perception rather than knowledge about their exact visual form.

But how is this possible? If word knowledge specifies what letters co-occur in words, then it seems to follow that word knowledge must be applied to the results of letter perception. But at the same time, the results of Reicher's experiment seem to show that letter perception is facilitated by word context. It seems like we are trying to have it both ways.

The paradox can be resolved if we observe that the process of letter identification need not be complete, exhaustive, or accessible to report before partial results of letter processing begin to interact with our knowledge about words (Massaro, 1975; McClelland, 1976). Perhaps what word knowledge does is allow us to reinforce partial, preattentive activations of letters that are consistent

with words in our vocabulary (Rumelhart, 1977). Suppose the presentation of a display produces activations of letters consistent with the display and that these in turn produce activations of words consistent with the letters. These activations, in turn, produce feedback to the letters, reinforcing the activations of letters consistent with the activated words (Adams, 1979). An illustration of this conception is presented in Fig. 2.1.

A second important clue to the process of word perception comes from the finding that the stimulus need not in fact even be a word to produce facilitation compared to unrelated letters or single letters. A large number of studies using Reicher's procedure have shown that the word advantage over unrelated letters extends to pronounceable nonwords as well as words (e.g., Baron & Thurston, 1973; Spoehr & Smith, 1975), and one study has obtained a large and reliable
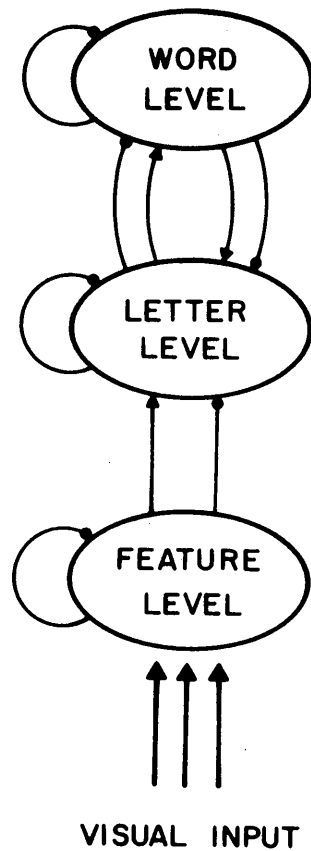

FIG.2.1   Sketch of the interactive activation model of word perception.

advantage of pseudowords compared to single letters (McClelland & Johnston, 1977).

One view that has often been taken concerning these findings holds that knowledge specific to words is not directly used in word perception. Instead, it is assumed that perceptual processing relies on a set of orthographic or phonological rules or on a parsing process that constructs representations of words and pseudowords in like manner (Spoehr & Smith, 1975), prior to a postperceptual lexical access process. It appears, however, that such a mechanism could not be the whole story, since in many experiments an advantage for words over pronounceable nonwords is obtained (Manelis, 1974; McClelland, 1976), even if it is not always statistically reliable (MeClelland & Johnston, 1977; Spoehr & Smith, 1975). Although attempts have been made to explain these differences away on the basis of possible orthographic and/or pronounceability differences between words and nonwords (Massaro, Venezky, & Taylor, 1979), the word–pseudoword difference does raise the possibility that word knowledge plays a sizable role in the perceptual processing of words. On the basis of this finding, McClelland and Johnston (1977) argued for two mechanisms—one in which familiar words were recognized, and a second in which representations of pronounceable strings are constructed. Similarly, Adams (1979) suggested that the words and nonwords both had an advantage over unrelated letter strings because of excitatory interactions between frequently co-occurring letters, whereas words had an advantage over regular nonwords because of feedback, as already mentioned.

Would it be possible to get by with a model that makes use of knowledge of specific words only? Recent work on the process of constructing pronunciations of pseudowords (Glushko, Chap. 3, this volume) suggested to us that it might. Glushko has shown that when constructing pronunciations of both words and nonwords, our knowledge of the pronunciations of specific words similar to the target word seems to influence the time and accuracy of our responding. On the basis of these findings, he has suggested that pronunciations of novel pseudowords are synthesized out of activations of pronunciations of the words that are similar to the target. In other words, we know how to pronounce REAT because we know how to pronounce words with similar spellings. These words, if partially activated in the process of perceiving REAT, would then produce partial activation of their corresponding phonological codes, which could then be synthesized to produce a pronunciation.

Much the same thing could be happening in the perception of pseudowords. Perhaps the presentation of a string like REAT produces partial activations of words with closely related spellings (for example, HEAT, REST, and READ, among others). These partial activations could then produce feedback to letter activations, just as in the case where an actual word stimulus is shown.

Thus, instead of explaining the perception of words and pseudowords in terms of a single mechanism that relies on orthographic or phonological knowledge, it

might be possible to explain the perception of both words and pseudowords on the basis of a single mechanism that relies solely on knowledge of specific words. Feedback from activations of specific words could reinforce activations of letters that happened to spell words, as well as activations of letters in strings that appeared to be close in spelling to some words. If, however, the string is a sequence of random unrelated letters (e.g., OMTP), it may not be similar enough in spelling to any words to generate much feedback.

There is one more finding that has played a central role in the development of our thinking. This is the fact that the perceptual advantage of words over non-words is greatly affected by the details of the visual conditions used. In particular, two sorts of visual conditions have been widely studied. In one, the target display itself is either very dim or very brief or both, and is followed either by no mask or a simple light mask. These conditions characterize most of the early work (pre-1970) on word perception, as well as a few more recent studies (e.g., Juola, Leavitt, & Choe, 1974; Rumelhart & Siple, 1974). Under these conditions, it seems reasonable to imagine that performance is limited largely by the quality of the information that can be extracted from the visual display. Large advantages for words over nonwords are obtained in these experiments, but only when a free-report measure of performance is used (Rumelhart & Siple, 1974; Smith, 1969); similar results are obtained for items presented in peripheral vision (Bouwhuis, 1979). When a forced-choice task is used, as in the Reicher procedure, there is only a slight advantage for single letters compared to letters in words, and the word advantage over nonwords is quite small (Johnston & McClelland, 1973; Juola, Leavitt, & Choe, 1974; Massaro & Klitzke, 1979). In contrast, when the target display is a bright, clear, high-contrast presentation of the word but is followed quite quickly by a high-contrast patterned mask, large advantages for words and pseudowords compared to single-letter and unrelated letter stimuli are observed.

A variety of different interpretations have been offered for the dependence of the word advantage on masking. One possibility is that the fact that the stimulus is a word makes it possible to maintain activations of a representation capturing the information in the display longer than would otherwise be the case, thereby increasing the chance that the subject would have sufficient time to translate the activated representation into a form suitable for overt report. The feedback mechanism we have been describing might have just such an effect. That is, feedback from activations of words could keep the representations of the letters active longer in the face of masking than would otherwise be the case. In no-mask conditions, this feedback would tend to reinforce possible letters that are consistent with the words that the subject knows but would tend to reinforce all the letters consistent with both the subject's knowledge of words and the visual information that he or she has managed to extract from the display. Qualitatively, then, we would expect the subject to pick a set of letters that form a word or a pronounceable nonword if asked to give a whole report. But he or she

should not show much of an advantage in discriminating between word (or pronounceable nonword) possibilities left open by information successfully extracted from the stimulus.

To summarize, the major findings we have reviewed thus far seem to be compatible with a model in which partial, preattentive activations of letters give rise to activations of words, which in turn produce feedback reinforcing the letter activations. We now review the model we have worked out to achieve this result. More detail is available in McClelland and Rumelhart (1980).

THE MODEL

We have already shown the general conception of the model in schematic form in Fig. 2.1. Perception is assumed to consist of a series of interacting levels, each level communicating with those immediately above and below it. In general, of course, a given level may have more than one level immediately above or below it, but for simplicity we now consider the case in which there is a linear ordering of levels. We have assumed that communication proceeds through a spreading activation mechanism in which activation at one level "spreads" to neighboring levels. Furthermore, we have assumed that the communication can consist of both excitatory and inhibitory messages. Excitatory messages increase the activation level of their recipients, and inhibitory messages decrease the activation level of their recipients. The arrows in the diagram represent excitatory connections, and the circular ends of the connections represent inhibitory connections. The intralevel inhibitory relationships represent a kind of lateral inhibitory relationship in which certain units at the same level compete. Thus, for example, since a string of four letters can be interpreted as, at most, one four-letter word, the various possible words mutually inhibit one another and in that way compete as possible interpretations of the string.

Although we assume that there are many levels that might be important in reading and perception in general and that the interactions among these levels are important for many phenomena, we have found that we can account for many of the major phenomena in word perception by considering only the interactions between "letter-level" and "word-level" elements. Thus, we have elaborated the model only on these two levels, have assumed that the other levels merely generate input into these levels, and have ignored the feedback that may occur between word and letter levels and any other levels of the system.

Specific Assumptions

For every relevant unit in the system there is an entity called a *node*. There is a node for each word, and there is a node for each letter in each position. The word nodes are located at the word level, and the letter nodes are located at the letter

level. Each node has connections to a number of other nodes. The set of nodes to which a node connects are called its *neighbors*. Each connection is two-way and may be either excitatory or inhibitory.

Connections occur both within levels and between adjacent levels. Connections within the "word level" are mutually inhibitory since only one word can occur at any one place at any one time. Connections between the word level and letter level may be either inhibitory or excitatory depending on whether or not the letter is a part of the word in the appropriate letter position. The set of nodes with excitatory connections to a given node are its excitatory neighbors, whereas the set of nodes with inhibitory connections to a given node are its inhibitory neighbors. A subset of the neighbors of the letter *t* are illustrated in Fig. 2.2.

Each node has a momentary level of activation associated with it. This level of activation is a real number bounded between a maximum or ceiling level and a minimum or floor level. Any node with a positive degree of activation is said to be *active*. In the absence of inputs from its neighbors, all nodes are assumed to decay back to an inactive state—that is, to an activation value at or below zero. This resting level may differ from node to node and corresponds to a kind of a
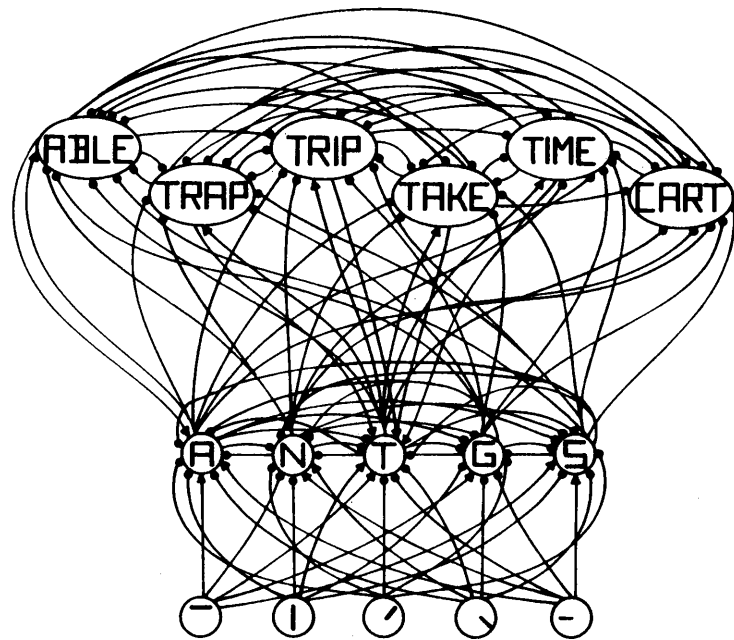

FIG. 2.2 A few of the neighbors of the letter *t*, in the first position in a word, and some of their interconnections. Excitatory connections are represented by arrows ending with points, and inhibitory connections are represented by arrows ending with dots.

priori bias (Broadbent, 1967), which might be affected by frequency of activation of the node over the long term.

When the neighbors of a node are active, they influence the activation of the node by either excitation or inhibition, depending on their relation to the node. These excitatory and inhibitory influences combine by a simple weighted average to yield a net input to the unit, which may be either excitatory (greater than 0) or inhibitory. Note that inactive nodes have no influence on their neighbors; only nodes in an active state have any effects, either excitatory or inhibitory.

The net input to a node drives the activation of the node up or down depending on whether it is positive or negative. The degree of the effect of the input on the node is modulated by the node's current activity level to keep the input to the node from driving it beyond its maximum and minimum values (Grossberg, 1978). The new value of the activation of a node at a given instant in time is equal to the value at the previous instant, minus the decay, plus the influence determined by the activations of its neighbors at the previous instant in time.

Upon presentation of a stimulus, a set of featural inputs are assumed to be made available to the system. For simplicity, the model assumes that the input consist of letters written in the font used by Rumelhart and Siple (1974) and shown in Fig. 2.3. During each moment in time, each feature has some probability less than or equal to 1 of being detected if it has not been detected already. Upon being detected, the feature begins sending activation to all letter-level nodes that contain that feature. All letter-level nodes that do not contain the extracted feature are inhibited. The probability of detection and the rate at which

FIG. 2.3. The features used to construct the letters and the letters themselves (from Rumelhart & Siple, 1974).

the feature excites or inhibits the relevant letter nodes are assumed to depend on the clarity of the visual display.

Presentation of a new display following an old one results in the probabilistic extraction of the set of features present in the new display. These features, when extracted, replace the old ones in corresponding positions. Thus, the presentation of an *O* following the *R* described earlier would result in the replacement of detected features already described with their opposites.

The Operation of the Model

Now consider what happens when an input reaches the system. Assume that all prior inputs have had an opportunity to decay, that the entire system is in its quiescent state, and that each node is at its resting level. The presentation of a stimulus initiates a chain in which certain features are extracted and excitatory and inhibitory pressures begin to act upon the letter-level nodes. The activation levels of certain of the letter nodes are pushed above their resting levels. Others receive predominantly inhibitory inputs and are pushed below their resting levels. These letter nodes, in turn, will begin to send activation to those word-level nodes of which they are a part and inhibit those word nodes with which they are not consistent. In addition, the various letter-level nodes attempt to suppress each other, with the strongest ones getting the upper hand. As word-level nodes become active, they in turn compete with one another and send excitation and inhibition back down to the letter-level nodes. If the input features are close to those for one particular set of letters and those letters are consistent with those forming a particular word, the positive feedback in the system will work to converge rapidly on the appropriate letter set and word. If not, they will compete with each other; and perhaps no single letter set nor single word will get enough activation to dominate the others, and their inhibitory relationships might strangle each other. The details of the process are greatly affected by the values of various parameters of the model. Some of these effects are described in McClelland and Rumelhart (1980).

Simulations

Though the basic ideas of the model are simple, its behavior is quite complex and cannot be understood without actually "running" it. We have been able to do this by simulating the performance of the model using a computer. To do this, we have had to make several simplifying assumptions. First, the simulation of the model operates in discrete time slices or ticks, updating the activations of all the nodes in the system once each cycle, on the basis of the values on the previous cycle. We have endeavored to keep the time slices "thin" enough so that the model's behavior is continuous for all intents and purposes.

Second, we have assumed that the weight of the excitatory and inhibitory effects of one node on another depend only on the levels at which the nodes are located. In other words, the strengths of the connections between all letter nodes and all word nodes of which they are part are assumed to be the same, independent of the identity of the words.

Finally, these simulations have been restricted to four-letter words. We have equipped our simulation with knowledge of 1179 four-letter words occurring at least two times per million in the Kucera and Francis word count (1967). Plurals, inflected forms, and occasional unfamiliar entries arising from apparent sampling flukes, acronyms, abbreviations, and proper nouns have been excluded.

An Example. For the purposes of this example, imagine that the visual display illustrated in Fig. 2.4 has been presented dimly to a subject on a CRT. In the first letter position, the letter *W* has been dimly presented. In the second letter position, the letter *O* has been dimly presented. In the third position, the letter *R* has been dimly presented. In the final position, only those features consistent with both the letter *K* and the letter *R* have been presented, and there is a blotch or something obscuring the features that would distinguish between these two possibilities. We wish now to chart the activity of the system resulting from this presentation. Figure 2.5 shows the time course of the activations for selected nodes at the word and letter levels.

At the word level, we have charted the activity levels of the nodes for the words *WORK, WORD, WEAR,* and *WEAK.* Note first that the word *WORK* is the only word in the. lexicon consistent with the presented information. As a result, its activation level is the highest and reaches a value of .8 through the first 40 time cycles. (The maximum and minimum activation values are set at 1.0 and −.2, respectively.) The word *WORD* is consistent with the bulk of the informa-
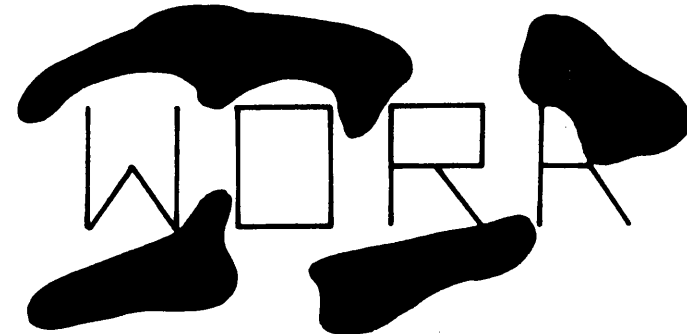

FIG. 2.4.   A hypothetical display presented to a subject in an experiment on word perception.

word activations

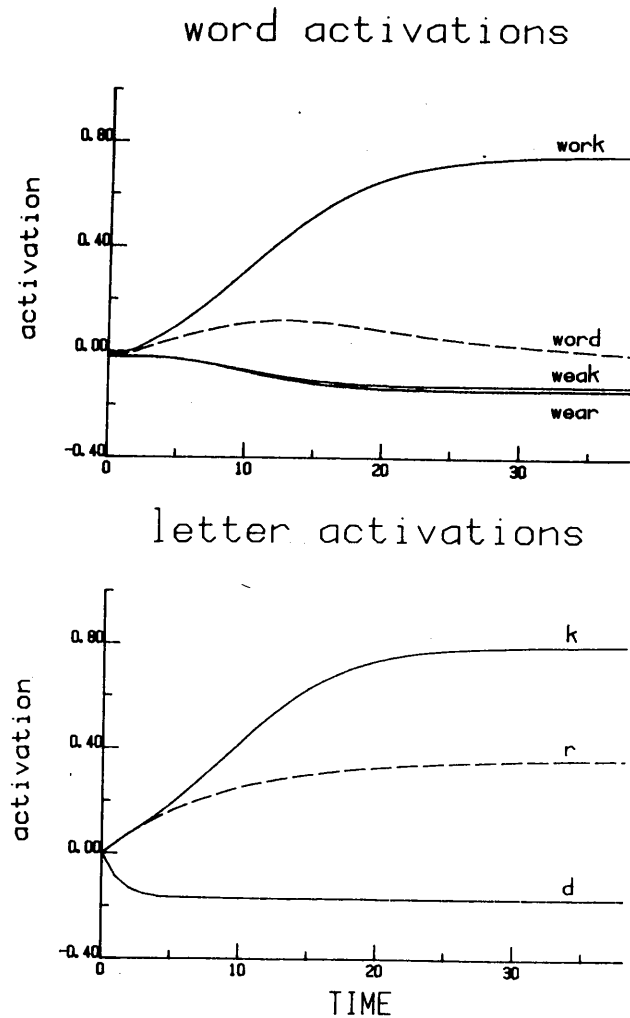

letter activations

TIME

FIG. 2.5.   The time course of activations of selected nodes at the word level (upper panel) and letter level (lower panel) after onset of the display shown in Fig. 2.4.

tion presented and, as a result, first rises and later—as a result of competition with *WORK*—is pushed back down to (and below) its resting level. The words *WEAR* and *WEAK* are consistent with the information presented in the first and fourth letter positions, but inconsistent with the information in the second and third letter positions. Thus, the activation levels of these words drop to a rather low level, but not as low as the activation levels of words such as *GILL*, which

contain *nothing* in common with the presented information. Although not shown in the figure, these words attain activation levels of about $-.20$ rather quickly and retain them throughout. Note that *WEAR* and *WEAK* are equally consistent with the presented information and thus drop together for the first nine or so time units. At this point, however, top-down information has determined that the correct final letter is *K* and not *R*. As a result, the word *WEAK* "becomes" more similar to the input than the word *WEAR* and, as a result, begins to gain a slight advantage over *WEAR*. This result occurs in the model because as the word *WORK* gains in activation level, it feeds back down to the letter level to strengthen the *K* differentially over the *R*. This strengthened *K* and weakened *R* in turn feed back into the word level and strengthen words ending in *K* and weaken those ending in *R*.

Now consider the activity at the letter level. Here we have plotted the activation levels of the letters *D, R,* and *K*. Note that at the start, the information clearly disconfirmed the existence of a *D* in the fourth position, and thus, the activation level of the *D* node decreased quickly to near its minimum value. However, the bottom-up information from the feature level supported both a *K* and an *R* in the fourth position. Thus, the activation level for each of these nodes rose slowly. These activation levels, along with those for *W, O,* and *R*, pushed the activation level of *WORK* above zero, and it began to feed back; by about Time Cycle 4, it was beginning to push the *K* above the *R* (*WORR* is not a word). Note that this separation occurred just before the words *WEAK* and *WEAR* separate. It is this feedback that causes them to separate. Ultimately, the *R* reaches a level well below that of *K*, where it remains, and the *K* pushes toward a .8 activation level. Note that in our simulations, we have adopted the simplifying assumption that there is no word-to-letter inhibition and no intraletter inhibition. Thus, *K* and *R* both coexist at moderately high levels—the *R* fed only from the bottom up, and the *K* fed from both the bottom up and the top down.

This example shows how our model permits relatively weak and ambiguous bottom-up information to be reinforced and enhanced by top-down processes. Here we have a very simple mechanism capable of parallel cooperative processing of words.

On Making Responses

One of the more problematic aspects of a model such as this one is a specification of how these relatively complex patterns of activity might be related to the content of percepts and the sorts of response probabilities we observe in experiments. The model assumes that the percept corresponds to a temporal integration of the pattern of activation over all the nodes. The integration process is assumed to occur slowly enough that very brief activations may come and go without necessarily entering perceptual experience or being accessible for purposes of responding; the longer an activation lasts, the more likely it is to be reportable.
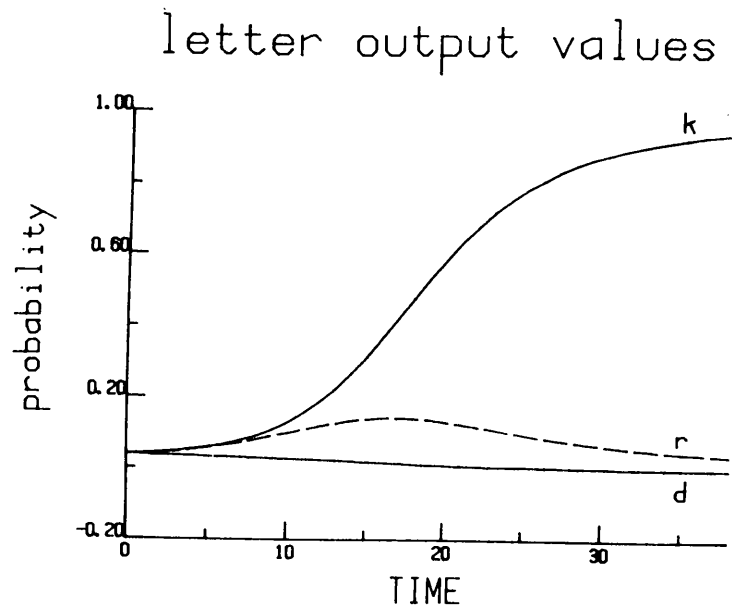
letter output values

FIG. 2.6.   ''Output values'' for the letters R, K, and D after presentation of the display shown in Fig. 2.4.

Specifically, we think of the response strength in the sense of Luce's choice model (Luce, 1959)—as being an exponential function of the activation of the relevant node averaged over the immediately preceding time interval. Following Luce's formulation, we assume that the probability of a response dependent on a particular node at a given level is equal to the ratio of the strength of that node, divided by the sum of the strengths of all other relevant nodes (e.g., nodes for letters in the same letter positions).

With regard to the previous example, it is useful to look at the "output values" for the letter nodes R, K, and D. Figure 2.6 shows the output values for these simulations. The output value for a particular letter at a particular time is the probability that the letter would be selected as the output or response if the response were selected at the given time. As intended, these output values grow somewhat more slowly than the values of the letter activations themselves, but eventually come to reflect the activations of the letter nodes as they reach and hold their asymptotic values.

APPLICATIONS OF THE MODEL TO THE LITERATURE

The model we have described turns out to provide a good account for a wide range of phenomena in the literature on word perception (McClelland, 1980). In this chapter, we lack the space to consider all the experiments that the model can

account for in detail. What we do instead is try to illustrate a few of the major features of the model's operation, in accounting for some of the most important facts.

The Word Advantage and Its Dependence on Masking

As mentioned previously, it is commonly observed that subjects are able to report the letters in a dim, degraded, or parafoveal presentation of a word more accurately than letters in unpronounceable nonwords. However, there is very little advantage for letters in words compared to nonwords under these conditions if Reicher's forced-choice test is used. Under these conditions, it is assumed that performance is limited quite simply by the quality of information that the subject has extracted from the display, and the information extraction process is not affected by feedback from the word level. What is affected is the probability of choosing particular letters that are compatible with the extracted visual information.

When the target is presented clearly but followed by a postdisplay mask, a large advantage for forced-choice performance on letters in words over single letters or letters in unpronounceable nonwords is obtained. Under these conditions, we assume that the target presentation results in complete feature extraction in all conditions, and performance is limited by the time available for encoding the activations produced by the presentation of the target before they are wiped out by the mask. The role of the word level is to provide feedback activations to the letter level that have the effect of increasing the strength of the activations of the correct letters at the letter level, thereby increasing the probability of correct encoding.

Figure 2.7 illustrates the behavior of the model in two trials of an experiment under this type of visual conditions. The target is the word READ in one case and the single letter E in the other. When READ is shown, the nodes for the letter E and all the other letters begin to get active soon after target presentation. These activations in turn activate the node for READ, which produces feedback, reinforcing the activation in the case where the E is presented in the context of the word. The mask drives down the activations quickly, but the area under the activation function is considerably larger in the case of the E in READ than in the case of the E alone. The result is that the probability of encoding the correct letter reaches a higher peak for the letters in READ than for the letter E alone. Note that as time goes on, the probability of encoding the correct letter drops back to chance. Thus, it is necessary to invoke the encoding process at the right time in order to ensure optimal performance. In applying the model to data, we assume that the subject learns when to ''read out'' as a result of practice in the task.

Perception of Pronounceable Nonwords

The mechanism just described is capable of producing facilitation for letters in nonwords as well as in words so long as the nonwords are similar to several
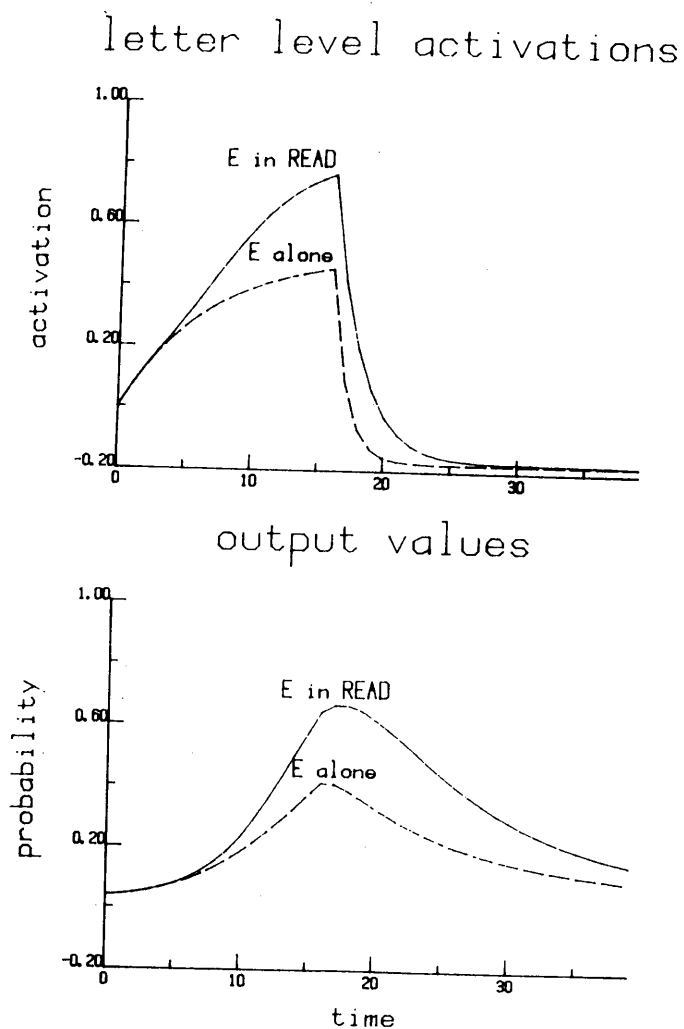
letter level activations

output values

time

FIG. 2.7.   The upper curve shows activation functions for the letter E in isolation and in the word *READ* under conditions of distinct target presentation followed by a distinct posttarget mask. The lower panel shows the output functions for the corresponding conditions.

words. The reason is that the presentation of a nonword similar to words results in activation of the nodes for the similar words, and these nodes then produce feedback, reinforcing the activation of the nodes for the letters. Of course, no actual word will reinforce the nodes for all the letters. However, if several words share one subset of letters in common and several others share other subsets, then

all the letters in the word can receive facilitation. For example, consider the pronounceable nonword *MAVE*. This item shares its last three letters with nine words, the first and the last two letters with one word, and the first two and the last letter with eight more words. All these words receive partial activation from the presentation of *MAVE*, and therefore, all the letters in *MAVE* receive some feedback reinforcing their activation. For this reason, the letters in *MAVE* will have a greater chance of being correctly encoded than the same letters in isolation.

There are now two papers in the literature showing that perception of letters in pronounceable nonwords is facilitated—compared to perception of letters in unrelated letter strings—only if the subject is aware that the list of stimuli may contain pronounceable nonwords. Interestingly enough, the facilitation for letters in words does not depend on expectation in this way. So, for example, if subjects expect nonwords, there is an advantage for letters in words, but not for letters in pronounceable nonwords. In fact, if they expect only words, there is an advantage for letters in words over pronounceable and unpronounceable nonwords and no advantage for pronounceable nonwords (Carr, Davidson, & Hawkins, 1978).

At first glance these results seem to suggest that there is one mechanism for the perception of words that is automatic, and a different mechanism for the perception of pronounceable nonwords that is only brought into play when pronounceable nonwords are expected. However, our model provides a very different type of account. We assume that the result depends on the ratio of excitation to inhibition in the influence of activations at the letter level on activations at the word level. Consider, first, the case where the inhibition is three times as strong as the excitation. In such a case, the presentation of four letters that have three letters in common with some word will have no effect on the perceptibility of the word. For example, presentation of *MAVE* will not result in the activation of any words. The reason is that the excitatory effect of the three letters that are compatible with a particular word would be canceled by the inhibitory effect of the fourth letter. So with high letter-word inhibition (relative to excitation), *MAVE* would produce no activations at the word level and no feedback-reinforcing activations at the letter level. But if the item is actually a familiar word (say, *CAVE*), the node for *CAVE* will receive strong activation because all four letters would excite the node and none would inhibit it. Thus, high letter-word inhibition will result in facilitation for letters in words but not for letters in nonwords.

Now consider again the case in which the letter-word inhibition is low relative to letter-word excitation. This is actually the case we considered initially when considering the partial activations of words that might result from the presentation of the nonword *MAVE*. Of course, such partial activations will also occur in cases where a word is actually shown. However, in this case, the node for the word shown still gets considerably more excitation than the nodes for other words. And because there is competition among nodes at the word level, the

nodes for words with only three letters in common with the word shown are kept from becoming strongly active. Thus, the node for the word shown is the major determinant of the amount of feedback whether the letter-to-word inhibition is low or high.

The Flexibility of the Perceptual System

One of the intriguing aspects of the literature on word perception has been that different results are often obtained in experiments that are superficially quite similar. What we see in the experiments already discussed is evidence that both the visual conditions of stimulus presentation and the expectations of the subject for what types of materials will be shown are important determinants of perceptual performance. The model nicely captures these effects. The model, like the human subject, behaves in different ways under different conditions, where the conditions include both the characteristics of the inputs to the system and the settings of various parameters in the system itself.

Other Effects in the Literature

There are a number of other phenomena in the literature that we have been able to account for with our model. One of these comes from an examination of the effects of intraword constraints on forced-choice performance by Johnston (1978). Johnston found that the letter *S* is no more easily seen in *SHIP* than in *SINK*, even though the context in the first case is consistent with (that is, forms a word with) only 3 possible letters, whereas the context in the latter case is consistent with 12 to 14, depending on what is counted as a word. Johnston used only words in the experiment, so we would expect that the subjects would have adopted a high value of letter–word inhibition. The visual conditions were like the distinct target/patterned mask conditions in which the word advantage over single letters is obtained. Under these conditions, only the nodes for the letters actually shown receive net excitation on the basis of stimulus input, and only the node for the single word containing all the active letters receives net excitation from the output of the active letters. Under these circumstances, the number of words quite similar to the target is irrelevant, because none of the corresponding nodes receive any net excitation. Thus, under the distinct target/patterned mask conditions of the Johnston study, the model produces no constraint effect either.

When a lower value of letter–word inhibition is adopted, the matter becomes more complicated. Under these circumstances, words that share three letters in common with the target do receive some activation, and as a result they compete with the target word. The more such words there are, the less active the node for the target word becomes, and therefore the less feedback there is to reinforce activations at the letter level. Thus, highly constraining contexts, in which few words share the same three context letters with the target, produce less word-

level interference than weakly constraining contexts, and therefore these items have a very slight advantage (a 1% difference in a simulation of accuracy differences in forced-choice responding). The reason is twofold. First, the node for the letter shown dominates the activation at the word level, and the weak activations of competing words have rather little effect. Second, there tend to be several words that share three letters with the target word including the critical letter. For example, in the case of *SINK*, there are words like *SANK*, *SILK*, *SING*, etc. Such words, when they become partially active, actually reinforce the activation of the critical letter and thereby help to dilute the effects of the words that have the three noncritical letters in common with the target.

Constraints do make a difference in the model when the input is degraded so that multiple-letter possibilities are partially activated by the input features. This is especially true if the response the subject must make is to identify what word is shown. Given incomplete feature input, relatively unique words—words sharing three letters in common with few other words—stand a greater chance of being uniquely compatible with the extracted features than words that have three letters in common with a large number of other words. For words that are themselves frequent, uniqueness is not so important, since such words tend to dominate the response of the system if other words are partially activated anyway. But for infrequent words, uniqueness is quite important, since it increases the likelihood that the correct node will be the one that is most compatible with the features extracted. When the node for a word that is more frequent than the word shown is as compatible with the extracted features as the actual word shown, the word shown tends to lose out to the more frequent word. Thus, the model has the following property: With degraded input, uniqueness is important for low-frequency words but not for high-frequency words. This characteristic of the model has been observed experimentally by Broadbent and Gregory (1968).

Another recent finding (Johnston & McClelland, in press) is that the word advantage over single letters is greatly reduced when the patterned mask contains letters rather than non-letter-patterned elements. Johnston & McClelland (in press) found a large word advantage with nonletter masks, but a much smaller advantage when the mask was made of letters. It did not matter whether or not the letters spelled a word.

In our model, this result comes about only in cases where the input is strong and of very high quality, so that the effect of feedback from the word level is to increase the persistence rather than the height of the activation function. The benefit for letters in words under these circumstances is due to the longer time available for readout of the activations at the letter level, when feedback causes them to persist for a longer time. The reason why the presence of letters in the mask reduces the word advantage is that the letters in the mask quickly produce activations of their own, which interfere with the readout process (recall that the probability of correct readout depends on the strength of the correct node, divided by the sums of the strengths of the other nodes). Letters in the mask only

have a very slight effect on performance of single-letter trials because the activations of the target letters are already wiped out by the mask by the time the activations of nodes for the letters in the mask have a chance to affect readout. In essence, the idea is that in the case of words, there is still something left for the mask letters to interfere with, but this is not so (or, rather, is so to a very limited extent) in the case of single letters.

In order to get the model to behave this way, we found it was necessary to restrict the inhibitory effect of the mask on active letter representations while making the excitatory effect of the mask letters on the nodes for the letters in the mask quite strong. It was only in this way that we were able to keep the mask from wiping out the old activations long enough to let the new ones due to the letters in the mask have an effect.

Context Enhancement Effects

In addition to accounting for all the findings already discussed, we have applied the model to a new set of findings that we call "context enhancement effects" (Rumelhart & McClelland, 1980). We briefly describe these effects, then consider how the model can account for them.

The basic idea of these studies is to use Reicher's forced-choice procedure to study how the information that subjects are given about the context of a target letter influences the identification of the target. To do this, the durations of the context and of the target letters are varied separately. For example, in our first experiment, the context could be turned on at various times—before, simultaneously with, or after the target letter. In this experiment, the target and mask letters were turned off together and replaced with a patterned mask. Figure 2.8 shows the results of this first study. What we see is that the accuracy of performance in identifying the target letter depends on the duration of the context. When the context is on 1.67 times as long as the target, performance in identifying the target is 15% more accurate than when the context is on only .33 times as long as the target, even though the target is on for the same amount of time in all cases. Thus, the longer the context information is available to the subject, the more accurately he or she can perceive a target letter in that context. Of course, since we used the Reicher procedure, it is not possible to attribute the effect simply to an influence of the context in helping the subject guess the target letter.

Our model has little difficulty in accounting for these results. The reason is that the longer the context is on, the more it contributes to the activation of the node for the whole word and the more this node, in turn, produces feedback reinforcing the correct alternative. Of course, if the context is on before the target letter, it tends to produce feedback exciting the node for the incorrect alternative as well, but because of the strong bottom-up inhibition from the feature to the letter level, the presentation of the target letter quickly nullifies such effects.

In a long series of experiments, we discovered several other important facts about these contextual enhancement effects. First, we demonstrated that the
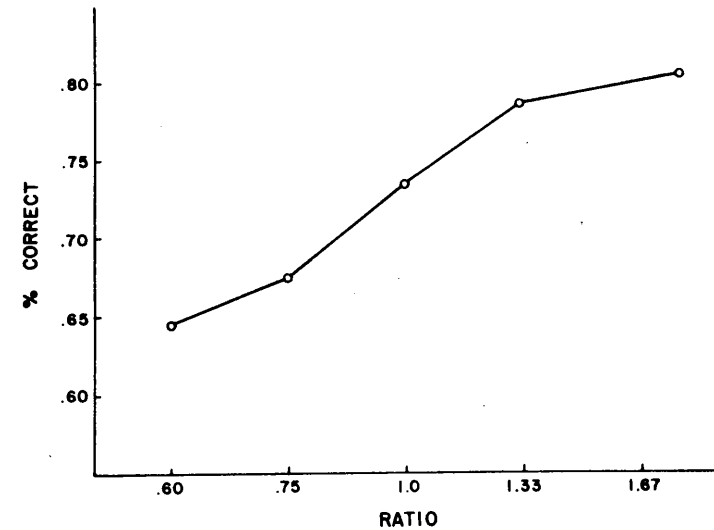
FIG. 2.8. Effect of context duration on probability correct forced choice on a critical target letter.

effect could not be obtained if the context consisted of three randomly chosen digits instead of the letters of a word. Second, we demonstrated that the effect was stronger if the extra context information preceded rather than followed the target letter. The effect of presenting the context early versus late was particularly strong for letters in the middle positions. Third, we carried out a complete factorial experiment, independently varying the duration of each of the four letters and the position of the letter tested. In this study, we were able to determine how much extra information about the letter in one position tended to facilitate perception of the letter in each other position. In general, the result was that the effect of extra information was greatest when it occurred in a position adjacent to the target letter.

The model accounts for all these effects, although additional assumptions are necessary to account for the fact that the internal letters benefit more from context than the external ones. The fact that completely unrelated context fails to produce facilitation is expected because completely unrelated context would fail to activate any higher-order nodes containing the target and the context, and thus would not afford any increase in feedback (or, for that matter, any feedback at all). The fact that the context must come earlier to be of much use also falls out of the simulations. Indeed, the model tends to overpredict the magnitude of this effect somewhat. The reason is that under the conditions assumed to prevail in simulating these results, the mask drives the activations of the letter nodes down so fast that the additional feedback to be gained from increased persistence of the activation of the relevant word node is of no benefit. As already noted, the model

does not directly account for the greater dependence of the interior letters on order of context presentation. One way to account for this is to assume that subjects ordinarily follow an outside-in processing strategy—first establishing the end letters and then using the information derived from them to facilitate perception of the letters in the middle of the item via feedback. Finally, it turns out that the greater effect of adjacent letters on perception of the target falls out of the model without any further assumptions. The reason has to do with the statistical properties of the language. It turns out that adjacent letters in a given word tend to co-occur in the same word more frequently than nonadjacent letters. The phonological structure of the language puts more constraints on what letters can occur right next to each other than it places on what letters can occur at a greater distance from each other in words. Following an initial P, for example, only a vowel or a liquid consonant (R or L) may occur. But the letter separated from P by another letter might be almost any letter at all. This fact about the language is latent in our model, since all the words are stored in the set of nodes. Thus, when a word like *PLAY* is presented, increasing the duration of P is likely to activate more words containing L than words containing A or Y.

The model, then, fares rather well in accounting for these findings, as well as several others that we have not had space enough to mention. However, the model does have some difficulty accounting for the magnitude of the context enhancement effect obtained with pronounceable nonwords. The problem may be seen easily by considering an example, such as *BINT*, with a forced choice between B and W, say. What happens when the context is presented, in the model, is that the _INT tend to activate all those words containing these three letters, such as *DINT, HINT, LINT*. Any activation they may provide for words that begin with a B will have to be weaker, since because the item is a nonword, the remaining letter does not make a word with all three of the context letters. These words containing B may be weakly activated, but they have to contend with the inhibition generated by the activations of the "enemies" of the initial B. For other cases, the model shows a substantial pseudoword enhancement effect. Overall, however, the model underestimates the degree of the effect. We have considered a number of modifications to our present model that can account for this underestimation, including the possible involvement of nodes for letter clusters, phonological codes, and nodes for words of more or less than four letters. For details of these modifications, see Rumelhart and McClelland (1980).

SUMMARY AND CONCLUSION

Our model is basically very simple, and it is clear that a complete model of the role of context in the perception of letters would be considerably more complex. Nevertheless, the interactive nature of the processes involved give the model remarkable power and flexibility. Yet the processes carried out by the model are

clearly not beyond the capability of simple neural circuits that we already know exist in the brain and nervous system. Through the use of computer simulation procedures we have been able to generate specific predictions from our model for nearly all the important phenomena of word perception. For the most part, these predictions have compared very favorably with observed results in a wide range of experimental studies. Obviously, however, there are many factors that go into reading and word perception that are not yet included in our model. Some of these factors are considered in some detail in McClelland and Rumelhart (1980) and in Rumelhart and McClelland (1980). Other factors have not yet been added to the model.

In addition to the substantive claims of the model we have developed, we are equally interested in the methodological implications of our work. We are becoming increasingly convinced that information-processing systems are properly viewed as consisting of numerous, relatively simple processing units (like our nodes) whose complexity results from the pattern of interactions that occur among the units and whose control structure is characterized by the modulation of levels of activation associated with the processing units. Such models, although simple in conception, gain their enormous processing power through the operation of populations of such units. We believe that computer simulations offer the clearest handle on the operations of such systems. Thus, we consider our efforts here as the beginning of an exploration into a kind of processing system seldom applied to phenomena as complex as reading. We are encouraged with our discoveries thus far.

ACKNOWLEDGMENTS

REFERENCES

Adams, M. J. Models of word recognition. *Cognitive Psychology*, 1979, *11*, 133-176.

Baron, J., & Thurston, I. An analysis of the word-superiority effect. *Cognitive Psychology*, 1973, *4*, 207-228.

Bouwhuis, D. G. *Visual recognition of words*. Eindhoven: Greve Offset B. V., 1979.

Broadbent, D. E. Word-frequency effect and response bias. *Psychological Review*, 1967, *74*, 1-15.

Broadbent, D. E., & Gregory, M. Visual perception of words differing in letter digram frequency. *Journal of Verbal Learning and Verbal Behavior*, 1968, *7*, 569-571.

Carr, T. H., Davidson, B. J., & Hawkins, H. L. Perceptual flexibility in word recognition: Strategies affect orthographic computation but not lexical access. *Journal of Experimental Psychology: Human Perception and Performance*, 1978, *4*, 674-690.

Cattell, J. M. The time taken up by cerebral operations. *Mind*, 1886, *11*, 220-242.

Grossberg, S. A theory of visual coding, memory, and development. In E. L. J. Leeuwenberg & H. F. Buffart (Eds.), *Formal theories of visual perception*. New York: Wiley, 1978.

Howes, D. H., & Solomon, R. L. Visual duration threshold as a function of word-probability. *Journal of Experimental Psychology*, 1951, *41*, 401-410.

Huey, E. B. *The psychology and pedagogy of reading*. New York: Macmillan, 1908.

Johnston, J. C. A test of the sophisticated guessing theory of word perception. *Cognitive Psychology*, 1978, *10*, 123-154.

Johnston, J. C., & McClelland, J. L. Visual factors in word perception. *Perception & Psychophysics*, 1973, *14*, 365-370.

Johnston, J. C., & McClelland, J. L. Experimental tests of a hierarchical model of word identification. *Journal of Verbal Learning and Verbal Behavior*, in press.

Juola, J. F., Leavitt, D. D., & Choe, C. S. Letter identification in word, nonword, and single letter displays. *Bulletin of the Psychonomic Society*, 1974, *4*, 278-280.

Kucera, H., & Francis, W. *Computational analysis of present-day American English*. Providence, R.I.: Brown University Press, 1967.

Luce, R. D. *Individual choice behavior*. New York: Wiley, 1959.

Manelis, L. The effect of meaningfulness in tachistoscopic word perception. *Perception & Psychophysics*, 1974, *16*, 182-192.

Massaro, D. W. (Ed.). *Understanding language: An information processing analysis of speech perception, reading, and psycholinguistics*. New York: Academic Press, 1975.

Massaro, D. W., & Klitzke, D. The role of lateral masking and orthographic structure in letter and word recognition. *Acta Psychologica*, 1979, *43*, 413-426.

Massaro, D. W., Venezky, R. L., & Taylor, G. A. Orthographic regularity, positional frequency, and visual processing of letter strings. *Journal of Experimental Psychology: General*, 1979, *108*, 107-124.

McClelland, J. Preliminary letter identification in the perception of words and nonwords. *Journal of Experimental Psychology: Human Perception and Performance*, 1976, *1*, 80-91.

McClelland, J., & Johnston, J. The role of familiar units in perception of words and nonwords. *Perception & Psychophysics*, 1977, *22*, 249-261.

McClelland, J. L., & Rumelhart, D. E. *An interactive activation model of the effect of context in perception, part I* (Chip Report 91). La Jolla, Calif.: Center for Human Information Processing, 1980.

Miller, G. A., Bruner, J. S., & Postman, L. Familiarity of letter sequences and tachistoscopic identification. *Journal of Genetic Psychology*, 1954, *50*, 129-139.

Rumelhart, D. E. Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Rumelhart, D. E., & McClelland, J. L. *An interactive activation model of the effect of context in perception, part II* (Chip Report 95). La Jolla, University of California, San Diego, Calif.: Center for Human Information Processing, 1980.

Rumelhart, D. E., & Siple, P. The process of recognizing tachistoscopically presented words. *Psychological Review*, 1974, *81*, 99-118.

Smith, F. Familiarity of configuration vs. discriminability of features in the visual identification of words. *Psychonomic Science*, 1969, *14*, 261-262.

Spoehr, K., & Smith, E. The role of orthographic and phonotactic rules in perceiving letter patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 1975, *1*, 21-34.