

An Interactive Activation Model of Context Effects in Letter Perception: Part 2. The Contextual Enhancement Effect and Some Tests and Extensions of the Model

David E. Rumelhart and James L. McClelland

An Interactive Activation Model of Context Effects in Letter Perception: Part 2. The Contextual Enhancement Effect and Some Tests and Extensions of the Model

David E. Rumelhart and James L. McClelland
University of California, San Diego

The interactive activation model of context effects in letter perception is reviewed, elaborated, and tested. According to the model context aids the perception of target letters as they are processed in the perceptual system. The implication that the duration and timing of the context in which a letter occurs should greatly influence the perceptibility of the target is confirmed by a series of experiments demonstrating that early or enhanced presentations of word and pronounceable-pseudoword contexts greatly increase the perceptibility of target letters. Also according to the model, letters in strings that share several letters with words should be equally perceptible whether they are orthographically regular and pronounceable (SLET) or irregular (SLNT) and should be much more perceptible than letters in contexts that share few letters with any word (XLQJ). This prediction is tested and confirmed. The basic results of all the experiments are accounted for, with some modification of parameters, although there are some discrepancies in detail. Several recent findings that seem to challenge the model are considered and a number of extensions are proposed.

Issues surrounding the role of familiarity and context in perception have been studied using stimuli consisting of letters since the beginning of experimental psychology. These studies show clearly that perceptual processes are affected by context and familiarity. In previous work one of us proposed an interactive model of reading to account for these and related effects (Rumelhart, 1977). The central feature of this model is that the processing of information in reading is assumed to consist of a series of levels. Information flows both directly, at once—from lower to higher levels and higher to lower levels. The proposal that information from a higher level can feed back and affect processing at a lower level explains how knowledge of a higher level unit, such

as a word, can affect the processing of a lower level unit, such as a letter.

In Part 1 of this paper (McClelland & Rumelhart, 1981), we combined the fundamental features of the Rumelhart (1977) interactive model with the flow-of-activation assumptions of the McClelland (1979) cascade model to build a new model called the interactive activation model. The model is capable of accounting for the fundamental facts of word perception, as verified by computer simulation of the results of a number of experiments demonstrating basic effects in the literature. The form of the model is illustrated in Figure 1.

In the model processing is organized into several levels. For simplicity we have limited our consideration to the three levels illustrated in the figure: the feature level, the letter level, and the word level. Each level consists of a set of units or nodes, one for each possible element at that level. Thus the word level consists of a set of word nodes, and the letter level consists of a set of letter nodes, one for each letter in each position within a word. The feature level consists of a node for each possible feature at each letter position.

Preparation of this paper was supported by NSF grant BNS-76-15024 to D. E. Rumelhart, by NSF grant BNS-76-24062 to J. L. McClelland, and by the Office of Naval Research under contract N00014-79-C-0323. We would like to thank Don Norman, James Johnston, and members of the LNR perception research group for helpful discussions of much of the material covered in this paper. Requests for reprints should be sent to David E. Rumelhart or James L. McClelland, Department of Psychology, C-009, University of California, San Diego, La Jolla, CA 92093.

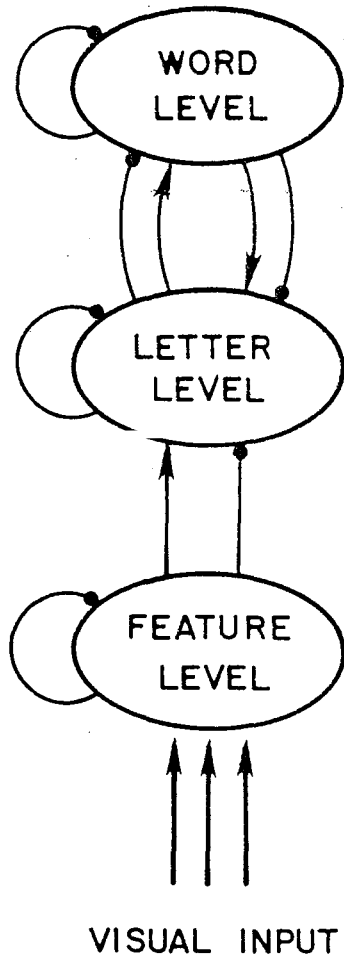


Figure 1. The various levels of processing considered in the interactive activation model and their interconnections. (Lines ending with arrows represent excitatory effects and those ending with dots indicate inhibitory interactions.)

activation level exceeds a threshold excites other nodes with which it is consistent (as, for example, an initial *T* is consistent with the word TAKE) and inhibits other nodes with which it is not consistent.

We assume that when a string of letters is presented to the visual system the feature level nodes are directly activated. Each feature node is assumed to activate all of those letter nodes consistent with it and inhibit all of those inconsistent with it. The more active a given feature node, the more it activates or inhibits the letter nodes to which it is connected. Inhibition and excitation are assumed to summate algebraically, and the net effect of the input on a node is modulated by the prior activation of that node. In this way those letter nodes with the most active feature nodes receive the most net excitation. At the letter level all nodes for letters in a given serial position are assumed to compete with one another through mutual inhibition. Each letter node is assumed to activate all of those word nodes consistent with it and inhibit all other word nodes. Each active word node competes with all the other word nodes and sends feedback excitation to the nodes for the letters consistent with it. Once a string of letters is presented and this process is set in motion the process will continue until either an asymptotic pattern of activation is reached, the input is turned off (and the activation of the individual units decay to a resting level), or a new stimulus (often a masking stimulus) is presented, thereby driving the system toward a new steady state and wiping out the remaining traces of the previous stimulus. In Part 1 of this article (McClelland & Rumelhart, 1981), we present a fuller description of the details of the model, and a discussion of the model's simulations of the basic findings in the existing literature on word perception.

In the present paper we elaborate and test the model, primarily against the results of previously unreported experiments. First we examine the role of the temporal relations between context and target-letter presentations in order to determine how well our model captures the actual temporal course of the facilitation that context provides for the perception of target letters. We will see that the duration and timing of the presen-

tation of the context with respect to the target letter greatly influences the perceptibility of the target letter, just as we would expect from a model like ours in which the perceptual facilitation of a letter depends on the ongoing processing of the letters in its context. We will see that the model is generally consistent with these effects, although some adjustment of parameters is required to make the model capture the beneficial effects of doubling the duration of the context when the display is a pronounceable pseudoword. We also consider two recent findings in the word perception literature that appear to challenge the model and discuss how the model may be consistent with these findings. Then we will test a counterintuitive prediction of the model—that unpronounceable and orthographically unacceptable nonwords made up entirely of consonants can produce as large a facilitation effect as pronounceable and orthographically acceptable items, if they share a number of letters in common with large numbers of words. Surprisingly, as we shall see, the prediction is supported by an experiment. Finally we suggest extensions of the model to three domains beyond the perception of letters in single tachistoscopic displays: the recognition of words in context, the pronunciation of visually presented words and pseudowords, and the perception of speech.

Temporal Relations Between Context and Target Displays

In the present model reading is treated as an interactive process in which contextual input is almost as important as direct evidence in the apprehension of stimulus material. The processing of a target letter in a multiletter display takes place within the context of the ongoing processing of the other letters, and processing of each letter is influenced by the effects of processing all of the others. For example, when a word is displayed, each letter helps activate the corresponding word node, and this node in turn strengthens the activations of each of the letter nodes. Thus as activation grows for one letter in a word it serves to facilitate the perception of the surrounding letters. It follows from this description of the perceptual

process that the duration and exact timing of the letters in a word context relative to the timing of a target letter should determine how much they can facilitate the perception of the target.

There has been little empirical investigation of the temporal relationship between target and related context. Estes (1975) has shown that presentation of context following the presentation of a target letter serves only to bias choices toward orthographically regular responses but has no effect on accuracy in a forced choice among orthographically similar alternatives. We know of no other published experiments that directly examine the effects of the temporal parameters of structurally related contexts on the perceptibility of the target letter. As we shall see our model makes various predictions about these temporal parameters, but there is no existing data base to test them against.¹ Therefore we undertook a series of experiments examining the effects of varying the temporal relations between target and context in words, pseudowords, and unstructured stimulus displays.

General Method

The method used in all of these studies had two main features: (a) We manipulated the onset and offset of each of the letters in the display separately, following the offset of each letter by a mask. (b) We tested the perceptibility of a single letter in the display on each trial, using Reicher's (1969) forced-choice test.

Figure 2 gives two ways the word WORK might be presented and illustrates the notational conventions we will use in discussing these experiments. Panel (a) illustrates a presentation in which the letters WOR_ are turned on at Time 0, followed a short time later by K. All letters are turned off simultaneously a little later. The K is the letter tested in the forced choice. Panel (b) illustrates a presentation in which the onset of the K precedes the onset of the WOR_.

¹ Johnston (Note 1) has done a number of similar experiments and obtained several findings similar to those reported here. This work was carried out concurrently with the present studies.

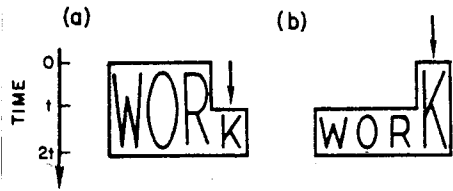


Figure 2. Notation used to represent durations of letters in different experimental conditions. (A display of target plus context is represented by a box. Time proceeds from top to bottom, letter position in the four-letter string from left to right. The arrow designates the target letter tested in the forced-choice test. Thus (a) indicates a condition in which the context letters were presented at Time 0 and left on until 2t. The target letter was turned on until 2t. In (b) the time relations of target and context are reversed. Though not indicated in the figure, a mask immediately follows the offset of each letter. In all experimental conditions the target letter occurs in all four letter positions. Thus, except when specially noted, the particular example of assignment of target and context letters is arbitrary.)

Again the *K* was probed, and again all letters were turned off simultaneously.

Experiment 1

Our first experiment examines whether the perceptibility of a letter in a word depends on the duration of the context letters in the display. Durations were adjusted by turning on the context letters at different

times relative to the onset of the target. All letters were turned off simultaneously and followed immediately by a mask.

From the model we would expect that the longer the duration of the context letters (that is, in this case, the earlier the onset with respect to the onset of the target), the more they will facilitate target perception.

Method

Procedure. The display conditions used in Experiment 1 are illustrated in Figure 3. The trial began with a fixation field. The subject pressed a button when ready and 250 msec later a string of four letters was presented. There were five different display-conditions characterized by the ratio of the duration of the context relative to the duration of the target. Ratios of .6, .75, 1.0, 1.33, and 1.67 were employed. Note that a ratio of 1.0 corresponds to a normal presentation in which the context and target letters are turned on and off simultaneously. In the 1.33 and 1.67 ratio conditions the onset of the context preceded the target letter, and in the .6 and .75 ratio conditions the onset of the target letter preceded the context. Trials for each ratio condition were mixed together in random order, and subjects were given no warning about which condition was coming up or which letter position contained the target. As illustrated in the figure, the mask consisted of overlapping X's and O's.

One hundred msec after the onset of the mask, a pair of letters appeared immediately above the target letter. The subject's task was to indicate which of the two letters had been presented in that position by pressing one of two buttons.

The experiment was controlled by a PDP9 computer. Stimuli were displayed on a CRT screen located about

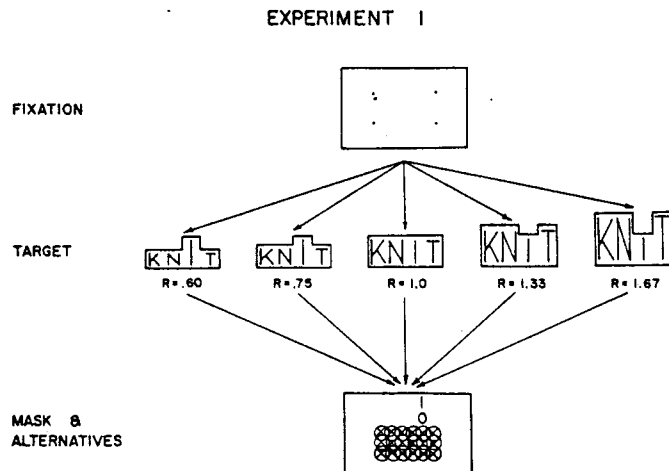


Figure 3. Display conditions used in Experiment 1.

40 cm in front of the subject. At this distance the four-letter display subtended a visual angle of about 30°. The duration of the target letter was adjusted for each subject after every block of 25 trials (including five of each type) to ensure an average of about 75% correct responses. An initial target duration of 50 msec for the target letter was employed. Thereafter the new duration was determined by the following equation:

$$d_{new} = d_{old} \left[1 + .75 \left(.75 - \frac{N}{25} \right) \right], \quad (1)$$

where *N* is the number of correct responses in the block. Duration of the context was based on the duration of the target.

Stimuli. The experiment was designed to study the perception of as many of the four-letter words in English as possible. The stimuli consisted of 1,500 words with frequencies of five or more per million (Kucera & Francis, 1967). The words were arranged into pairs differing in a single letter position. The frequencies of the two members of each pair differed by no more than a factor of two. Each member of each pair was seen by half of the subjects. The letter position in which a pair of words differed was the target position for that pair, and the letter in that position was the target letter. In the forced-choice test the alternatives were the target letter and the letter in the corresponding position in the other member of the word pair. All available pairs within the above constraints were used, so the number of tests in each serial position was not constant: 303 pairs differed in the first serial position, 118 pairs in the second, 153 in the third, and 176 in the fourth. (Complete lists of the stimuli used in Experiments 1 through 9 are available in Rumelhart & McClelland, Note 2.)

Subjects. Ten undergraduates at the University of California, San Diego were given either course credit or \$2 for serving in the experiment.

Results and Discussion

Subjects were unaware of the fact that the onsets of the different letters in the display varied. Phenomenologically some words were easier to see than others, but the onset-time differentials were small enough that all four letters seemed to come on and go off simultaneously. Nevertheless performance on the two-alternative forced choice was strongly affected by the quality of the context.

The results are shown in Figure 4. For the lowest ratios subjects responded correctly less than 65% of the time. For the highest ratios they responded correctly over 80% of the time. These points are based on a total of 1,500 observations in each condition, and the 95% confidence interval around each point is about ±2.5%.

Let us consider how our model would ac-

count for the fact that contexts turned on prior to the target produce more accurate perception of the target. When the context is turned on it begins to activate the nodes for the letters it contains. These letter nodes activate the nodes for words consistent with the context, including the nodes for the yet-to-be-presented target letter and the alternative. These word nodes strengthen the nodes for the context letters and awaken the nodes for letters completing the candidate words. Then, when the target letter is turned on, the letter strength can quickly grow and reach a relatively high value. The other primed letter-nodes are quickly inhibited due to the mismatch to the actual input. Figure 5 indicates the activations resulting from presentation of the word SHIP for the word nodes *ship* and *whip*, for the letter nodes *s* and *w*, and for the probabilities of selecting *s* and *w* as outputs for context to target ratios of 1:1 and 2:1. Clearly, presenting the context for twice as long as the target letter has the effect of increasing target selection and therefore forced-choice accuracy, just as we observed in the data.

In order to see if the magnitude of the effect produced by the model is about the same as that observed in the data, we ran a simulation of the experiment. In this and subsequent simulations, a sample of ten word-pairs differing in each serial position was chosen. One element of each pair was

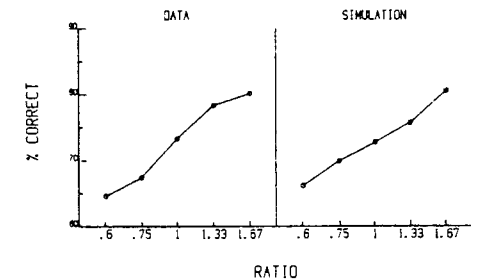


Figure 4. Percent of correct forced-choice responses to the target letter as a function of the relative duration of the context to the duration of the target letter. (The panel on the left shows the actual data from the experiment; the one on the right shows the results of the simulation run described in the text.)

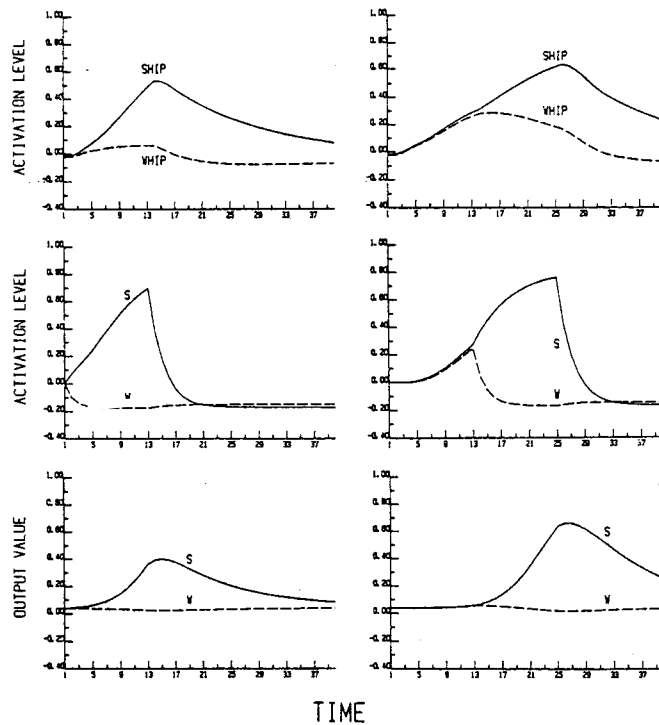


Figure 5. Activations and the word and letter levels and output values resulting from the presentation of the word SHIP, in the case where the S and the HIP are presented in a 1:1 and a 2:1 ratio.

chosen to be presented to the model. Thus, 40 items were used in the simulation. The parameters used here were those employed in the simulations reported in Part 1 (McClelland & Rumelhart, 1981). The duration used was 13 cycles. Optimal readout time was 3 cycles after the onset of the mask.

Figure 4 shows the results. The model provides a good account of the effects obtained empirically. For a ratio of .6 our simulation yielded about 66% correct. For a ratio of 1.67 the simulation yielded 81% correct. This is exactly the same range of values observed in our experiment.

Experiment 2

It may be argued that the effects of enhancing the context observed in Experiment

1 merely reflect some artifact of the peculiar timing sequences used. Perhaps, for example, the results are due to some sort of warning-signal effect rather than to the information that one letter can dynamically contribute to the processing of its neighbor letters. To test this hypothesis the effects of the enhancement manipulation were assessed on letters embedded in numerals, and on letters embedded in words.

Method

The procedural details were nearly identical to those of Experiment 1. Each of 10 subjects viewed 750 four-character displays. Context duration (2:1 and 1:1 ratios) and context type (word or numeral) were factorially manipulated within subjects. For each subject a different random assignment of 187 items to each of the word-context conditions and 188 items to each of the numeral

conditions was used. Numeral-context displays were generated by replacing the context letters from a word display with a set of three randomly chosen numerals.

Results and Discussion

The longer context duration enhanced target perception for letters in words, but no such enhancement was found with the numeral context (Figure 6). The enhancement effect was significant for letters in words, $F(1, 9) = 11.925, p < .01$, but not for letters in numerals, $F(1, 9) = .573, p > .5$, and the interaction of context with enhancement condition was significant, $F(1, 9) = 14.817, p < .005$. It appears, then, that the effect of context duration truly depends on the nature of the context. There is no indication that the context-enhancement effect is merely a warning-signal effect or some other artifact of this sort.

Our simulation of this experiment (also shown in Figure 6) was obtained with a duration of 15 cycles and a readout time of 2 cycles after the onset of the mask. The simulation produced the same basic effects as the actual experiment, though there are slight discrepancies. First, performance on letters in numeral contexts was about 5% worse in the actual data than in the simulation. It is possible that the number contexts were confusing the subjects in ways that

word contexts do not. If so this would lead to worse performance with numerals than with no context at all. Second, the model overpredicts the size of the enhancement effect for words. This overprediction appears to be due to the fact that the empirically obtained enhancement effect for words is reduced in size as overall word performance gets above 80% correct. In some of the experiments reported below, the enhancement effect for words is much larger when the overall performance level on words is lower. The model is not susceptible to whatever causes this reduction in the size of the effect at performance levels in the 80-90% correct range.

Experiment 3

If, as suggested by our model, the contextual information is having its effect while the target letter is being processed in the perceptual system, the exact timing of the extra contextual information should be very important. In our model when the contextual information comes on early it primes the node for the word shown, and this in turn primes the node of the target letter, thereby facilitating target perception. If the context followed the presentation of the target letter it should not help very much, because the mask quickly wipes out the activation produced by the target and leaves nothing for the context to facilitate. This experiment tests these implications of the model.

Method

The design of this experiment is illustrated in Figure 7. All letters were presented for the same duration, but the order of presentation varied. Three conditions were used in this experiment.

Context-early condition. In this condition the contextual information was presented first. The target was presented at the same instant that the context was turned off.

Simultaneous condition. In this condition context and target were turned on and off simultaneously.

Context-late condition. In this condition the target letter was presented first. The context was presented at the same instant that the target was turned off.

The offset of each letter was followed immediately by a mask for that letter. To allow for this the mask was changed from Experiment 1 to the one illustrated in Figure 7. Except for this change the display conditions of the experiment were exactly the same as those of Experiment 1.

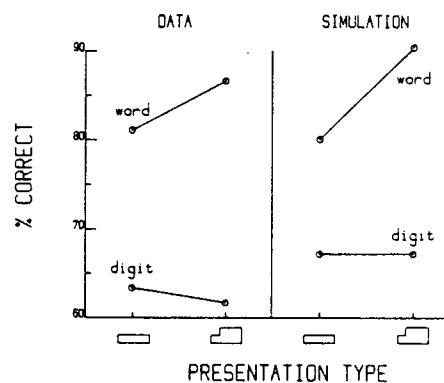


Figure 6. Effects of doubling the duration of the context on the perception of letters in words and in strings of numerals from Experiment 2. (Actual data are shown on the left; the results of the simulation are shown on the right.)

EXPERIMENT 3

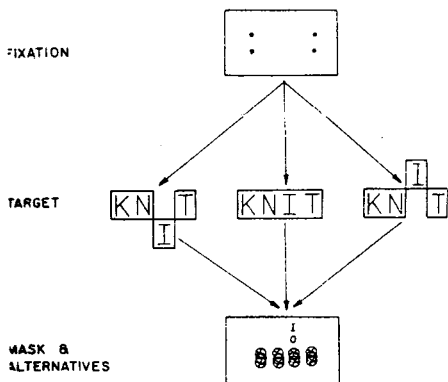


Figure 7. Illustration of the design of Experiment 3, investigating the effects of presenting the context before, after, or simultaneous with the target.

The procedure also adhered to the conventions established in Experiment 1. Each of 10 subjects viewed 240 context-late trials, 270 simultaneous trials, and 240 context-early trials.

Results and Discussion

As expected, perception of the target letter was much better when the context came first than when the target came first (Table 1). However the early context did not seem to produce superior overall perception compared to the simultaneous context.

Our simulation run for this experiment (also shown in Table 1) used a duration of 14 cycles with the readout occurring 2 cycles after the offset of the target letter. The simulation agrees well with the data on the relative inferiority of the context-late condition. However the simulation produces a 6% ad-

Table 1 Experiment 3: Proportion Correct Responses as a Function of the Relative Times of Offset for the Target and Context Letters

Result class	Presentation condition		
	Context late	Simultaneous	Context early
Observed	.682	.742	.749
Simulation	.651	.765	.827

vantage for the context-early condition compared to the simultaneous condition.

A clue to the reason for the discrepancy is given in Figure 8, which shows the serial-position curves for each of the conditions. The curves are relatively flat for the simultaneous condition but not for the other conditions. The context-late condition shows a U-shaped curve typical of random letter strings, and the context-early curve forms an inverted U. It appears that processing may somehow be proceeding from the outside in, so that outside letters have an advantage if they are presented early in the interval but have a disadvantage when they are presented late. Such a mechanism is obviously missing from our model. In a later section we will discuss how such a mechanism might be incorporated.

For the time being it is worth noting that when we look at the data for letters in the middle of the word we see a pattern very similar to that observed in our simulations (Table 2). Thus the difference between the overall simulation results and the overall results of the experiment appears to be entirely owing to effects on the first and last serial positions.

In this experiment we have been able to control the times at which the context information was available relative to the target and thereby manipulate the effect of the contextual information on the perceptibility of the target letter. The pattern of these effects would seem to confirm to a substantial de-

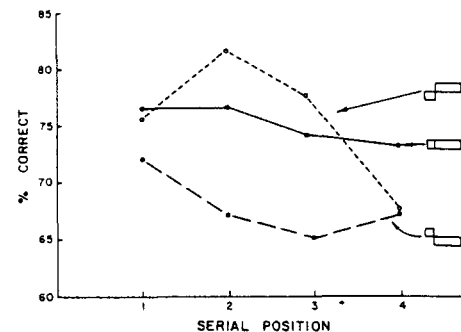


Figure 8. Serial-position curves for the context-early, simultaneous, and context-late conditions of Experiment 3.

Table 2 Experiment 3: Proportion Correct for Serial Positions 2 and 3 as a Function of Presentation Condition

Result class	Presentation condition					
	Serial position 2			Serial position 3		
	Context late	Normal	Context early	Context late	Normal	Context early
Observed	.671	.764	.813	.650	.741	.770
Simulated	.649	.774	.845	.653	.765	.791

gree the assumptions of the model concerning the priming effect of the context letters on the perceptibility of the target letters.

Experiment 4

Suppose that we leave the context information on for a fixed interval and simply vary the place in the interval when the target information is available. In our model if the target is presented early the extra contextual information would have less time to prime the relevant word nodes than if the target information were presented later in the interval. Experiment 4 tests this prediction. It also reintroduces the digit context to control for possible masking or warning-signal effects of the asynchronous presentation of target and context.

Method

Design. The design of this experiment is illustrated in Figure 9. The critical letter was presented for the same duration in all conditions, and the context was always presented twice as long as the target letter was. On half the trials the target occurred at the beginning of the interval defined by the onset and offset of the context, and on the other half of the trials it occurred at the end. Half of the time the context fit together with the target to make a word and half of the time the contextual letters were replaced by random numbers.

Stimuli. For Experiment 4 we used a new set of 384 word pairs with 96 pairs for each serial position. There was evidence from the previous experiments that performance was somewhat worse for words beginning with a vowel (whether the first serial position was tested), so all of the items in the list began with consonants.

Procedure. The procedure for this experiment differed from the procedure of the previously described experiments in a few details. The experiment was controlled by a PDP11 computer rather than the PDP9 computer used in the previous experiments. The fixation

point was modified from that illustrated in Figure 3 to the one illustrated in Figure 9. Trials were entirely self-paced. After the onset of the fixation point, subjects advanced to the presentation of the stimulus by pressing a button. The onset of the context display occurred 250 msec after the button was pressed. Each of 32 subjects (chosen as in Experiment 1) were given 384 trials, including one member of each pair. Across subjects each member of each pair occurred equally often.

Results and Discussion

As shown in Figure 10, target letters were much better perceived in word contexts than in numeral contexts. For numeral contexts there is no advantage when the target letter comes late, and in fact there is a very slight difference in the opposite direction. However with word stimuli there is a significant 4% advantage favoring the late target-letter condition, $F(1, 31) = 5.21, p < .05$. The interaction between context type and presentation condition is also significant, $F(1, 31) = 6.53, p < .05$.

The figure also shows our simulation results for this experiment. The simulation results show the same general pattern of results as those we have observed but with two discrepancies. First, performance in the numeral contexts is slightly worse in the data than in our simulation. Second, the obtained differences between early and late presentations are somewhat smaller than we obtained in the simulation. Possibly the model is overestimating the speed with which the mask affects the target letter. When the mask is turned on in our simulations, it immediately begins to reduce the activation of the target letter thus rendering totally ineffective subsequent contextual input. Another possibility, as mentioned before, is that

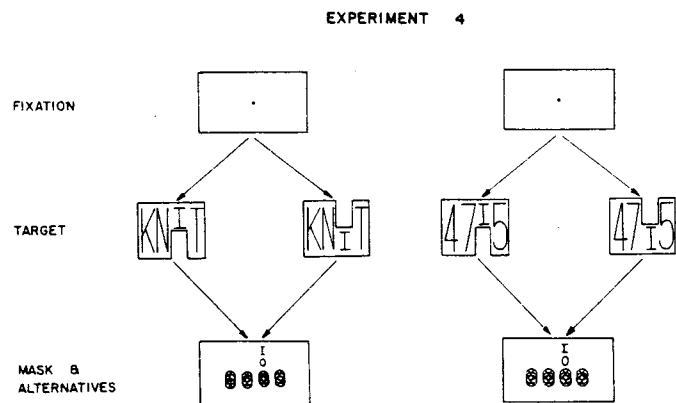


Figure 9. Illustration of the design of Experiment 4 in which the target letter is presented either early or late in the presentation period containing the context. (As in Experiment 2, the context either forms a word with the target letter or it is a random sequence of digits.)

the size of the effect is attenuated in the data because of the high overall performance level with words. In spite of these problems, the results of the experiment and of the simulations are basically consistent with the expectation that contextual information does prime those letters consistent with the con-

text and thereby aid in the perception of those letters.

Experiment 5

According to our model contextual information can substitute for a lack of direct sensory information, and conversely direct

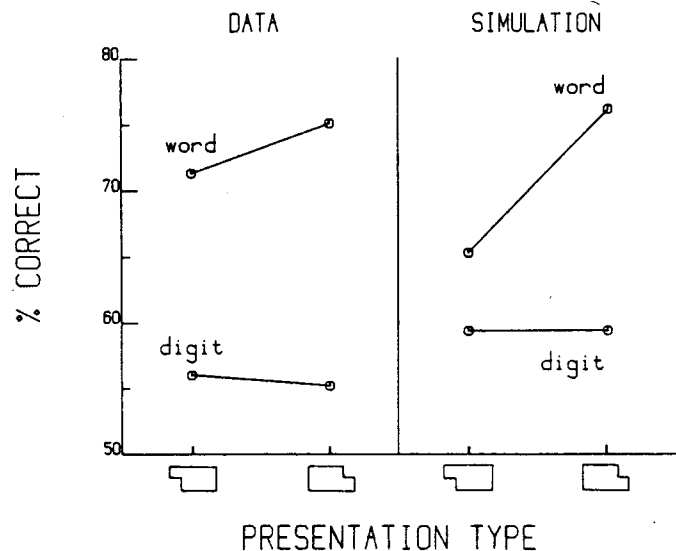


Figure 10. Percent correct responses for strings containing word and digit contexts as a function of whether the target letter came early or late in the interval. (From Experiment 4.)

sensory information can substitute for contextual information. In Experiment 5 we compared the relative contribution of additional direct evidence as compared to additional contextual information.

Method

The design of this experiment is illustrated in Figure 11. The target letter was presented for either duration D or duration $2D$, and the context was presented for either D or $2D$ msec, independent of the duration of the target letter. D was adjusted between blocks of trials to ensure a 75% correct response rate for each subject. Ten subjects were run using the stimuli and display conditions of Experiments 1-3.

Results and Discussion

The results shown in Figure 12 show the expected trade-off of direct and indirect information. Increasing the duration of either the target or the context increases performance substantially. The direct information is somewhat stronger than the indirect, but the effects appear to be additive. Both main effects are highly significant but the interaction is not. Our simulation also produced an additive effect of direct and indirect information. The effect of direct information is somewhat stronger in the simulation than in the observed data but otherwise data and experiment agree.

EXPERIMENT 5

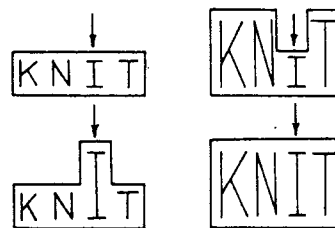


Figure 11. Illustration of the design of Experiment 5, comparing the effects of direct information (duration of the target letter) and indirect information (duration of the context letters).

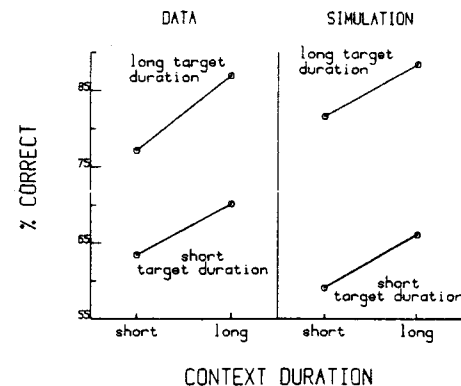


Figure 12. Percent correct responses as a function of the duration of the context and target presentations from Experiment 5. (Actual data are shown on the left; the results of the simulation are shown on the right.)

Experiment 6

We have been arguing throughout this section that the perceptibility of the target letter depends on the perceptibility of the letters in the context. Does this dependency depend on which context letters are enhanced? If so can our model account for such dependencies? The present experiment examines these issues.

Method

The conditions of Experiment 6 are illustrated in Figure 13. Each of the four serial positions was tested under eight context conditions. In each of these conditions, a different (possibly null) subset of the context letters was presented for twice as long as the target was. The onset of the enhanced context-letters always preceded the onset of the target, and all letters in the display were always turned off simultaneously. The stimulus set and the physical conditions were those of Experiments 1-3. Each of 24 subjects was run for 750 trials.

Results and Discussion

The results of both the experiment and our simulation are broken down condition by condition in Figure 13. The simulation produced about the same performance levels for all serial positions whereas the actual experiment did not. Overall, however, both simulation and experiment demonstrated

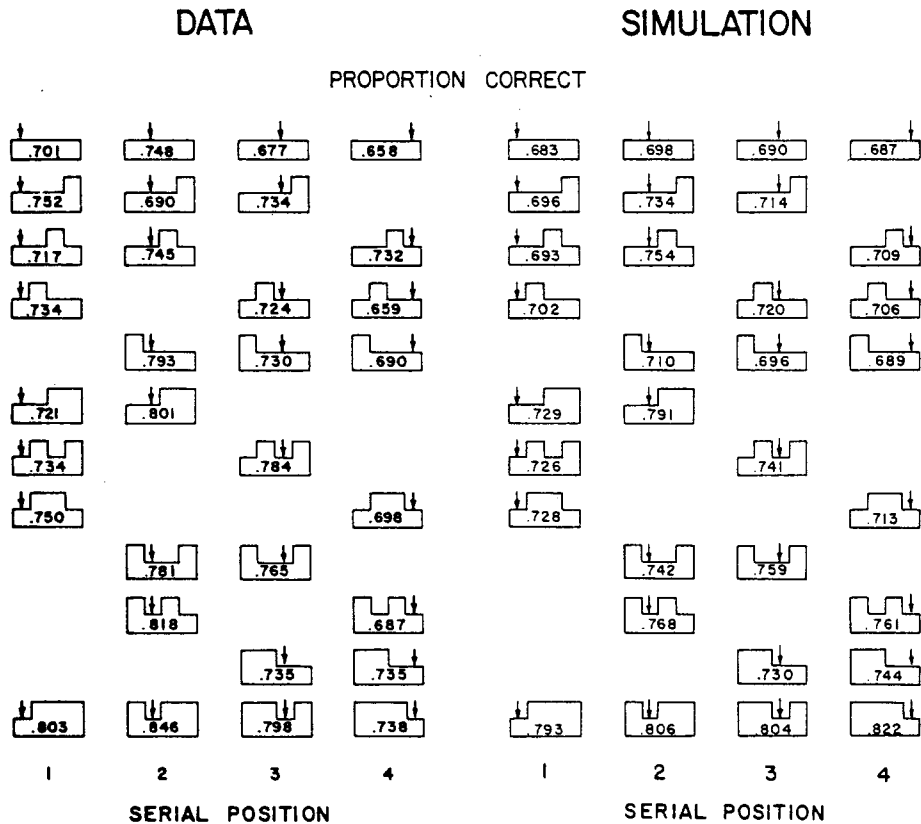


Figure 13. Illustration of the 32 different conditions used in Experiment 6 and percent correct (forced-choice) for each condition. (Actual data are on the left; the results of the simulation are shown on the right.)

that target accuracy increased with the number of context letters enhanced (Figure 14). Two other findings emerged from an analysis in which we determined the effect of increases in the duration of Letter l, on the detectability of Letter l. This effect was determined by computing the difference between the average percent correct on Letter l as a function of whether Letter l was enhanced. In the actual data (Table 3), the effect of context enhancement appears to be greater for adjacent letters than for separated letters. In addition, initial and final letters appear to have stronger effects than internal letters.

The simulation produced somewhat different results (Table 4). Adjacent context letters show a greater benefit for target letters in Serial Positions 3 and 4 but not for targets in Positions 1 and 2. Furthermore the generally stronger effect of end letters is not as evident here as it is in the actual data. This latter discrepancy is presumably related to the absence of performance differences as a function of serial position in the model. It is interesting to consider why the model shows any effects of the relative position of the target letter and enhanced context-letters. At first glance we would expect no such effects because the feedback is based on ac-

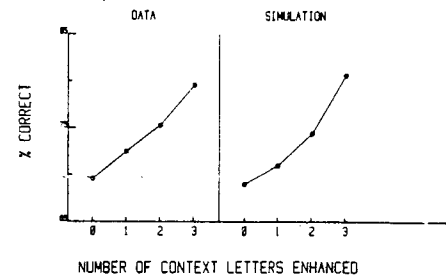


Figure 14. Average percent correct responses over all conditions as a function of the number of letters of context to receive the longer duration presentation.

tivation of nodes at the word level, and all four letter-positions feed activation to these nodes on an equal basis. However it turns out that, on the average, the closer two letters are in a word, the more knowledge of one tells us about the other. That is, the closer two letters are in a word, the more likely they are to occur together in many words in the language. Thus the "adjacency effect" exhibited by our model derives from the fact that nearby letters are more likely to activate words containing the target letter than are more distant letters. The failure of the adjacency effect to show up in all serial positions seems to be due to characteristics of the particular sample of items used in the simulation. The simulations were performed over a subset of 10 words for each serial position. It turns out that all of the first-position items in the simulation began with a single consonant followed by a vowel. In such items the likelihood of co-occurrence of the first and second letters tends not to be high because each consonant can occur with each

Table 4
Size of Context-Enhancement Effect as a Function of Serial Position of Enhanced Context-Letter: Simulated

Target position	Target position			
	1	2	3	4
1	—	1.3	3.1	4.3
2	3.7	—	3.4	4.2
3	3.4	5.9	—	5.2
4	3.5	3.6	4.6	—

vowel in words in English, with very few restrictions.

Experiment 7

In the experiments reported thus far we have investigated the effect of context enhancement on the perception of letters in words. These experiments can be thought of as extensions of Reicher's (1969) finding that presenting a letter in a word context presented at the same time as a target enhances perception of that letter compared to the presentation of the letter alone. Of course the fact that a string forms a word is not an essential characteristic for Reicher's effect; it can be obtained with pronounceable pseudowords as well as words, though not with unrelated-letter strings. Are analogous effects obtained when letters in such contexts are enhanced? The next three experiments constitute an investigation of the effects of context enhancement for pseudowords and other sorts of nonword strings. The first experiment of this series demonstrates that the context-enhancement effect does indeed occur with pronounceable pseudowords and shows that the size of this effect is comparable to that obtained with words.

Method

The procedure and the set of word stimuli used were those described in Experiment 4. One pair of pseudowords was formed for each pair of words by changing the letter most distant from the target letter to yield a pronounceable nonword. Vowels were replaced by vowels and consonants by consonants. In the case of the pair WORD-WARD, for example, we changed the final D and replaced it with an L yielding the pair WORL-WARL. This procedure ensures that the same target-letter pairs are tested in words and pseudowords and that the words

Table 3
Size of Context-Enhancement Effect as a Function of Serial Position of Enhanced Context-Letter: Observed

Target position	Target position			
	1	2	3	4
1	—	3.9	1.1	3.3
2	6.4	—	4.0	.4
3	2.7	3.3	—	5.3
4	2.6	1.6	2.9	—

and pseudowords are similar in consonant/vowel structure and have the same immediate context surrounding the target letter. The result of this procedure was a list of 384 stimulus quadruples of the form WORD-WARD-WORL-WARL.

Each of the 16 subjects viewed one member of each stimulus quadruple. Each subject saw 96 words and 96 pseudowords in the enhanced-context condition (2:1 context to target ratio), and 96 words and 96 pseudowords in the normal-context condition. Each item was tested equally often in each condition.

Results and Discussion

Though subjects were more accurate on words than pseudowords, a context-enhancement effect was obtained for both words and pseudowords (Figure 15). The overall effect of context enhancement was highly reliable, $F(1, 15) = 25.74, p < .001$, as was the effect of context type, $F(1, 15) = 23.74, p < .001$, but there was no interaction with context type, $F(1, 15) = 1.79, p > .1$, although the trend suggests that the enhancement effect may be slightly larger for words than for pseudowords.

In order to investigate the model's account of the enhancement effect with pseudowords, we simulated this experiment using the same sample of words used in all of our previous simulations. The pseudowords were the items that were actually paired with the sample words in the experiment. With the standard parameters we had been using up to this point, the model did not produce a pseudo-

word-enhancement effect. The simulation and the data showed the same pattern of results for the words and for the normal presentation of the pseudowords. The preview of the context, however, actually lead to slightly poorer performance on the pseudowords, averaging over all the items in the sample.

Clearly the behavior of our model is at variance with the facts. However it turns out that changes in two of the parameters were sufficient to bring the simulation back into line with the data. The changes were necessary to handle two problems that seemed to be keeping the pseudowords from showing an enhancement effect. The basic problem stems from the fact that with pseudowords there are no four-letter words containing all three of the context letters in the preview and the target letter. Words that match the advanced context perfectly do not contain the target letter—if they did the display would be a word—whereas words containing the target letter never match more than two of the three context letters. The result is that words that match the three context letters produce feedback that activates letter-level competitors to the target letter before it is even presented. These words also inhibit the words that contain the target letter and some of the context letters by lateral inhibition at the word level. This makes it more difficult for these words to exceed threshold later, when the target letter is actually presented. To avoid these interference effects, it appears to be necessary to suppose that some preactivation of word units can take place before they begin to produce feedback to the letter level and before they begin to inhibit each other. Thus the major change required was to adjust the resting levels of all words downward by .2, so that the resting levels of the highest frequency words were near -.2 and the resting levels of the lowest frequency words were near -.25. To accommodate this change the minimum activation value for word nodes was set to -.3. This change, in itself, was not sufficient to solve the problem of competition at the word level completely, particularly for items for which there exist high-frequency competitors consistent with all three context letters. In addition, it appeared to be necessary to keep

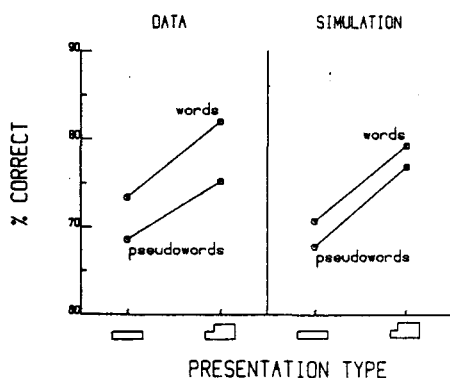


Figure 15. Percent of correct responses for words and pseudowords as a function of presentation type. (Actual data are shown on the left; the results of the simulation are shown on the right.)

active word nodes from inhibiting each other until their activations reached the value of .075. With these two changes plus minor tuning of one other parameter (the letter-to-word inhibition parameter was changed from .04 to .02), we were able to provide a reasonably close account of the word and pseudoword data from this experiment and the remaining experiments to be reported in the rest of this paper.

The results of the simulation for the present experiment using the altered parameters are shown in Figure 15. The simulation results are rather close to the actual experimental results. The major difference is that the overall difference between words and pseudowords is about 2% lower in the simulated results than in the experiment. The size of the enhancement effect, however, for both words and pseudowords is of the appropriate magnitude. Notice that there is no interaction between the presentation type and the word-pseudoword variable in either case.

The changes in the parameters did not affect the model's account for the results of the enhancement experiments described previously in which only words were used. Although reducing the resting level by .2 does delay the onset of feedback, the increase in the threshold for inhibition at the word level permits more words to participate in the feedback. Below we will consider the effects of these changes on the results of the simulations reported in Part 1 (McClelland & Rumelhart, 1981).

Experiment 8

The previous experiment shows that pronounceable nonwords show essentially the same pattern of interaction with context as words do. The present experiment examines whether the same is true for unpronounceable and orthographically irregular nonwords. According to the model the effect is due to partial activation of word nodes by word and pseudoword stimuli and thus should not be obtained with nonwords that are not similar to words.

Method

Experiment 8 was identical to Experiment 7 except that the nonword stimuli were constructed in a different

way. In this case the nonwords were made by a rearrangement of the letters of the word stimuli. For every word a nonword was constructed by reversing the order of the first and second letters and the third and fourth letters. Thus if the original order of the letters was 1-2-3-4, the new order would be 2-1-4-3. For four-letter words beginning with consonants this leads to an unusual string of letters, often unpronounceable. For example, the new quadruples containing the words WORD and WARD would contain the nonwords OWDR and AWDR.

Results and Discussion

The results from this experiment are shown in Figure 16. In order to maintain the same average percent correct, the durations had to be increased over those in the previous experiment with the resulting higher percent correct for the words. Even then the performance on the reversed words was somewhat poorer than for the nonwords of the previous experiment. Probably because of the restriction of the range, the enhancement effect on words was somewhat reduced in comparison to some of the previous experiments. Nevertheless it was highly significant, $F(1, 15) = 8.672, p < .001$. The enhancement effect for the reverse words, on the other hand, was much smaller and nonsignificant, $F(1, 15) = .747, p > .5$.

The simulation of the results of this experiment used the same word-pairs used in previous simulations. The pairs of nonwords were constructed by reversing the first two letters and the last two letters in the words.

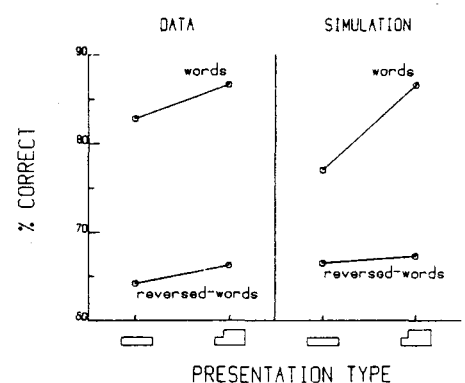


Figure 16. Percent of correct responses for words and reversed words as a function of presentation type from Experiment 8. (Actual data are shown on the left; the results of the simulation are shown on the right.)

The results of the simulation are illustrated in Figure 16. Once again the simulation does not appear to be susceptible to the ceiling effect that apparently reduced the magnitude of the enhancement effect on words in this experiment. Otherwise the simulated and actual results are comparable. In particular, neither the simulation nor the actual experiment produced much facilitation for reversed words.

Experiment 9

Whether an item is a word is a matter of fact, but whether it is a pseudoword is a matter of theory. According to several views whether an item is a pseudoword is a matter of degree, depending on how high the pseudoword ranks on probabilistic measures of string predictability based on one or more statistical constraints found in the words of English (Massaro, Venezky, & Taylor, 1979; Miller, Bruner, & Postman, 1954; Rumelhart & Siple, 1974). Indeed each of the cited articles provides evidence that probability of correct letter-identification and other performance measures do correlate with string predictability as calculated in various different ways. In the present experiment we examined whether predictability of letters in nonwords was correlated with performance in the forced-choice test, and whether predictability influenced the size of the context-enhancement effect. We also consider whether our model is consistent with such effects.

Method

Stimuli. The goal was to get a set of stimuli all of which were at least marginally pronounceable and orthographically regular but which differed in their predictability based on the statistical regularities of English. In order to do this a simple grammar of the four-letter words of English was constructed and the "set of possible four-letter words" of English was generated. Following this the actual words of English were culled by removal of all strings from the list that appeared in the Kucera and Francis (1967) word count and all other strings that were recognizable as words. The measure of predictability was the sum of the conditional probabilities of each letter given both the preceding and following context, according to the following equation:

$$V_i = p(L_1) + p(L_2|L_1) + p(L_3|C_1L_2) + p(L_4|C_1C_2L_3) + p(L_4) + p(L_3|L_4) + p(L_2|L_3C_4) + p(L_1|L_2C_3C_4). \quad (2)$$

L_j represents the j th letter of string V_i ; C_j represents the class (i.e., whether the letter was a consonant, a vowel, or a final E) of the j th letter of the string. The expression $p(A|B)$ represents the proportion of word types containing letter A in the appropriate position compared to those ending in pattern B . Items were counted as words only if they occurred at least five times per million in the Kucera-Francis word count. The equation gives a kind of summed measure of the extent to which we might be able to predict what each of the letters in the string might be based on each of the other letters.

Once values were assigned to all of the strings, the strings were ordered from best to worst. Quadruples of strings were then constructed so that there were two high predictability or "good" strings that differed by one letter and two low predictability or "poor" strings that differed by the same letters (e.g., BLAY-GLAY-BIPO-GIPO). A total of 384 such quadruples were generated, 96 for each of the four serial positions.

Procedure. The 16 subjects were run on a 2 x 2 design crossing string quality (good vs. poor) with context enhancement (2:1 context to target duration vs. normal context). Each subject saw one member of each stimulus quadruple, 96 in each condition.

Results and Discussion

The results for this experiment show that not all pseudowords are equally easy to see (Figure 17). The good pseudowords showed substantially better performance than the poor ones, though the difference is only marginally reliable, $F(1, 15) = 4.19, p < .07$. Both kinds of words seemed to show an improvement with a preview of the context, though the effect was substantially smaller for the poor strings than for the good ones.

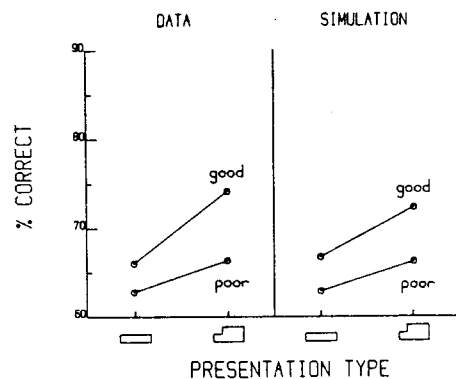


Figure 17. Percent of correct responses for good and poor pseudowords as a function of presentation type. (Actual data are shown on the left; the results of the simulation are shown on the right.)

The effect was highly significant for the good pseudowords, $F(1, 15) = 24.94, p < .001$, but only marginal for the poor items, $F(1, 15) = 4.30, p < .06$, and the interaction evident in Figure 17 was significant, $F(1, 15) = 5.70, p < .05$.

In our model the statistical constraints on letter predictability are not explicitly stored, yet the model performs better on strings that conform to these constraints just as the human observer does. In order to compare our model with the results of this experiment, we drew a sample of items from the new list of items used in the experiment. A total of 40 quadruples of words were chosen at random, 10 for each serial position. In the simulations we used one good and one poor pseudoword from each quadruple. Employing the same parameter values as in the two previous simulation runs, with a target duration of 12 time cycles, we obtained the results shown in Figure 17. A comparison with the actual data indicates that the model comes pretty close to the data. Clearly, the measure used to define the good and poor pseudowords is related to those factors that determine both the overall accuracy of performance and the magnitude of the enhancement effect in our model.

Serial-Position Effects

In several of the experiments we have reported, our model failed to account for the effects of serial position that were quite evident in the data. With normal presentations of word displays, performance varies little over serial position. However when the context is enhanced or temporally offset with respect to the target letter, strong serial-position effects emerge. These effects suggest that subjects use some sort of "outside-in" processing strategy that leads to variations in performance across serial position.

There are at least two possible ways of accounting for serial-position effects within the framework of our model.

1. The quality of the information at the ends of the words might be better than the quality of the information about letters internal to the word due to lateral interference (Eriksen & Rohrbaugh, 1970; Estes, All-

meyer, & Reder, 1976) or focus of attention. We can simulate the effect of varying stimulus quality or attention by adjusting the rate of activation of letter nodes. The idea is simply that the higher the quality of the input or the more attention devoted to it, the faster it should drive the relevant letter node toward its maximal activation level. In all of the simulations we have presented thus far, we have assumed fixed feature-to-letter influences, independent of serial position.

2. It may be that not all letters are read out simultaneously. In all of the simulation results reported, we have assumed that all letters are read out simultaneously at a time that results in optimal performance overall. It is, of course, possible that different serial positions are read out at different times, perhaps because the readout process demands limited resources.

In this section we show that implementation of these possibilities in our model allows us to account for some of the effects of serial position, and for their interaction with context-timing conditions.

We examine the effects of serial position for standard and enhanced conditions with word stimuli. It turns out that the major trends in these curves can be accounted for solely by supposing that the input rate is higher for some letters, particularly the first letter, than it is for others. The data used came from the standard and enhanced word-conditions of Experiment 7 and from like conditions run in another experiment, not described above, that used the same stimuli. The serial-position curves are illustrated in the left panel of Figure 18. These data show a bow-shaped serial-position curve under standard conditions and a relatively flat serial-position curve under enhanced conditions. The normal parameters used in the original simulation of Experiment 7 produce the serial-position curves of the central panel of the figure. However if we differentially weight the inputs to each of the four serial positions (giving the positions relative rate parameters of 1.6, 1.15, .85 and 1.05, respectively), we get the serial-position curves shown in the right panel of the figure. Clearly the results of the simulation capture the major features of the observed data. In-

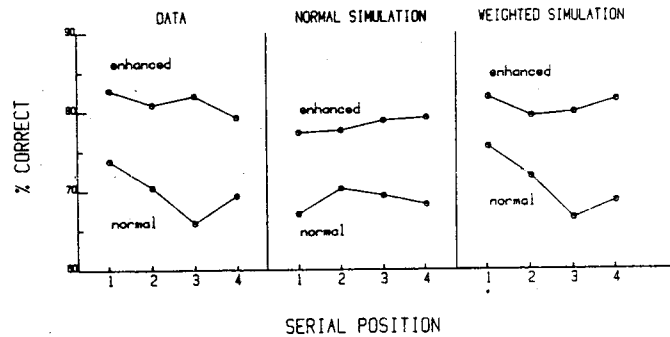


Figure 18. Interaction of serial-position and enhancement effects. (The left panel shows data from the word trials in Experiment 7 combined with data from similar conditions of an experiment not reported here. The center panel shows the results of a simulation run using standard parameters. The right panel shows the results of a simulation run in which the input strength varied across serial position.)

Interestingly the differential weights give the normally presented words the bow-shaped serial-position curve as we would expect while retaining the flat serial-position curve for the enhanced presentations. The reason for this appears to be that the perceptibility of the letters in the enhanced condition is more dependent on contextual information and less dependent on the direct information about that letter. The letters with the weakest direct input get the most help from the other letters.

The mechanism just described does not successfully account for the way the form of the serial-position curve varies as a function of the order of presentation of context and target as observed in Experiment 3. To account for these results, we combined the assumption of differential activation rate discussed above with the assumption that the readout occurred at different times for different positions. Specifically we assumed that the two end-letters are read out first, followed by the second letter three ticks later, and then the third letter three ticks later still. To optimize overall performance, readout for the end letters actually occurs two cycles before mask onset. This keeps readout for the third letter from occurring far too late. The results of this simulation are shown in Figure 19. A comparison with Figure 8 shows that we have captured the general features of the serial-position curves,

although the simulation produces much better performance in the fourth serial position for the context-early condition than we find in the actual data.

In addition to these two possible mechanisms, there are several other factors that might be contributing to serial-position effects. These include perceptibility differences of the particular letters that happen to occur in the different positions, statistical properties of the words with target letters in particular positions, variations in locus of fixation and attentional strategy as a func-

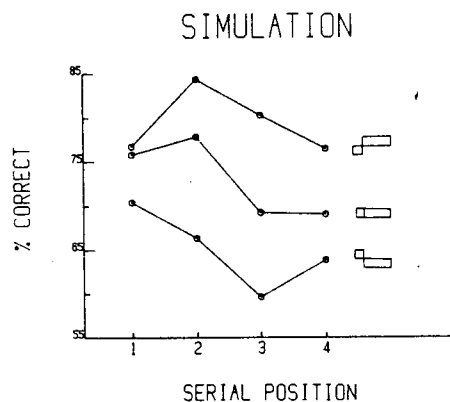


Figure 19. Simulated serial-position curves for the context-early, context-late, and simultaneous presentation conditions used in Experiment 3.

tion of experimental conditions and instructions, and so on. Because both of the mechanisms we have described are potentially subject to attentional control, it may be difficult to gain control and definitive understanding of these mechanisms until the factors that govern control of attention are understood. For these reasons we feel it may not be worthwhile to try to track down every aspect of the serial-position effects we have observed at this stage. The mechanisms we have suggested seem to be capable of accounting for the general trends in the serial-position data in a relatively straightforward manner, however, and so may be worth further consideration in later research.

Summary of Findings on the Effects of Context Enhancement

Our experiments have shown that the perceptibility of a target is strongly dependent on the duration and timing of the presentation of contextual information. The main finding was simply that the longer the duration of the context and the larger the number of context letters enhanced, the more accurate was forced-choice performance on the target letter. Enhanced context improved performance on the critical letter even though the context could not directly help the subject select the correct forced-choice alternative. Thus the context must affect perception of the target letter itself as it is being processed by the perceptual system. This conclusion is reinforced by the fact that the usefulness of the extra contextual information was increased when it came before rather than after the target letter was presented.

We also demonstrated effects of context enhancement with pseudowords as well as words. Clearly enhancement of contextual information helps letters in nonwords that are similar to words as well as actual words. However context enhancement has no effect on letters embedded in unrelated numerals or scrambled words and has only a slight effect on pseudowords that conform poorly to the statistical regularities of English letter strings.

The model we developed in Part 1 (McClelland & Rumelhart, 1981) to ac-

count for the effects of normally displayed contexts on letter perception also accounts for the effects of context enhancement. The experiments involving enhancement of pseudowords required a modification of the parameters, however.

Can the results of all of the experiments we have considered be accommodated with the revised parameters needed to account for the pseudoword-enhancement effect? To address this question we reran a sample of items from each of the experiments simulated in Part 1 and this part of the paper. We found that the model produced qualitatively equivalent results with one exception. The addition of the .075 threshold prior to interaction within the word level caused the model to generate a bigram-frequency effect of about 5% for pseudowords, whereas McClelland and Johnston (1977) found a negligible effect of this variable. The reason for the bigram-frequency effect in the model is that letters in high bigram-frequency pseudowords tend to occur more in words that partially match the item shown (see Part 1 for discussion). If all of the items are inhibiting each other, as with the previous parameter values, there is no net effect of bigram frequency on performance. The existence of more items is canceled out by more inhibition. The addition of the threshold before which items can feed back, but not inhibit each other, gives the advantage to the high bigram-frequency pseudowords.

It is not clear whether the discrepancy represents an inadequacy of the model. There may well be a slight bigram-frequency effect in pseudowords that went undetected in the McClelland and Johnston (1977) study. In fact, McClelland and Johnston actually found a slight effect in free reports of pseudowords, though it did not show up in the noisier forced-choice measure. Further, Experiment 9 produced a slight difference between our good and poor pseudowords, and the measure we used to categorize these items is highly correlated with bigram frequency.

In any case, aside from this discrepancy, our model is capable of accounting for all of the findings we have discussed thus far. Straightforward modifications seem called for to handle the effects of serial position.

There are, however, two recent findings in the literature that seem to support points of view on perceptual processing other than our own. We now consider these findings in turn.

Two Challenges to the Model

Effects of Set on Performance With Pseudowords

It appears that facilitation of the perception of letters in pseudowords does not occur unless the subject expects that pseudowords may be shown. Aderman and Smith (1971) found no reliable benefit of pseudoword context when subjects expected only unrelated letters. Carr, Davidson, and Hawkins (1978) replicated this result and added two more interesting facts (Table 5). First, they found that the word advantage over unrelated letters can be obtained when subjects expect only unrelated letters, even though letters in pseudowords show no reliable advantage under these conditions. Second, when subjects expect only words, they perform as poorly on letters in pseudowords as they do when they expect unrelated letters.

At first glance these data seem to suggest that there must be different processing mechanisms responsible for the word and pseudoword effects. There seems to be a word mechanism that is engaged automatically if the stimulus is a word and a pseudoword mechanism that is brought into play only if pseudowords are expected. However we will show that these results are completely consistent with our model, even though it has only a single mechanism for processing both words and pseudowords.

Let us recall how the model accounts for the pseudoword advantage in the first place.

Table 5
Effect of Expected Stimulus Type on the Word and Pseudoword Advantage Over Unrelated Letters (Difference in Probability Correct Forced Choice; Carr et al., 1978)

Target	Expectation		
	Word	Pseudo-word	Unrelated letters
Word	.15	.15	.16
Pseudoword	.03	.11	-.02

When four letters are presented, they activate the detectors for the presented letters. These, in turn, activate words that have two or more letters in common with the word shown. None of these words get strongly activated, but their aggregate activation is generally enough to reinforce the activations of the letters about as much as they would be reinforced if they formed an actual word.

Obviously activation of detectors for words that are not completely consistent with the four letters shown depends on the relative values of the letter-word excitation and inhibition parameters. If the inhibition is set to zero, the letters shown will tend to produce partial activations of all words that match any one or more of the active letters. Some of these activations will of course be squashed by lateral inhibition, but many will persist. As the inhibition increases it will tend to cancel the excitation, first for the words that match the input in only one position, then those that match in two, and finally those that match in three out of the four positions. Indeed if the letter-to-word inhibition is equal to three times the letter-to-word excitation, then no four-letter nonword can activate the node for any four-letter word. Even if the nonword has three letters in common with the word, the inhibition generated by the letter that is different will cancel the excitation generated by the letters that are the same. Thus as the letter-word inhibition increases, relative to the letter-word excitation, the extent to which the presentation of a pseudoword will tend to produce activations at the word level will decrease, vanishing when the letter-to-word inhibition reaches a value three times as great as the letter-to-word excitation. At that value the model produces no activations at the word level and therefore no advantage for letters in pseudowords over letters in unrelated-letter strings.

This argument suggests that we can account for the effects of set on performance with pseudowords by supposing that subjects control the letter-word inhibition parameter in our model. We need only assume that they use a low value when they expect pseudowords, and a high value when they do not. But we have still to consider what effects variation of letter-to-word inhibition might

have when the display actually spells a word. If relaxation of letter-to-word inhibition increases accuracy for letters in pseudowords, we might expect it to do the same thing for letters in words. However in general this is not the case. One factor to account for this is that the word shown still gets considerably more activation than any other word and tends to keep the activations of other nodes from getting very strong. A second factor is that activations of other words are not an unmixed blessing. These activations produce inhibition that keeps the activation of the node for the word shown from getting as strongly activated as it otherwise would. The third factor is that the activations of any one word sharing three letters with the word shown only reinforce three of the four letters in the display. For these reasons it turns out that the value of letter-to-word inhibition can vary from 0 to .21 with very little effect on word performance. Thus as the letter-to-word inhibition varies from 0 to 3 times the letter-to-word excitation, the model produces large variations in the size of the pseudoword advantage with no effect on the size of the advantage for words.

It does appear, then, that we can now account for Carr et al.'s (1978) findings by simply assuming that when subjects expect only words or only unrelated-letter strings they adopt a large value of the letter-to-word inhibition parameter, but when they expect pseudowords they adopt a small value. Perhaps a large value of letter-to-word inhibition is the normal setting, with a relaxation only if pronounceable pseudowords are known to be included in the list of stimuli. Generally speaking this strategy would appear to be a reasonable one. After all in the normal course of events in reading one is trying to read words rather than pseudowords, and it might be advantageous to keep words that are only similar to the word shown from becoming activated. On the other hand, if the item is not a word, then partial activations of words might be advantageous, not only to facilitate perception of the letters, but also to aid in the determination of a plausible pronunciation for the unfamiliar sequence of letters (Glushko, 1979).

Under conditions of degraded input, subjects would have to adopt a low value of let-

ter-to-word inhibition even if they expected to see words. A high value will not allow any words to become active when the input is sufficiently impoverished that there are several letter nodes in the same position that are partially activated on the basis of the available feature information. When multiple letter-nodes are active in the same position, each inhibits all of the words the others excite, and unless letter-to-word inhibition is weaker than letter-to-word excitation, it will only take two active alternatives in each position to keep any activity from occurring at the word level.

Effects of Masks Containing Letters

Recently, Johnston and McClelland (1980) have reported a series of experiments that support a hierarchical model of word perception in which there is no feedback and no within-level interactions among units. In the model there is bottom-up excitation and inhibition of letter detectors by feature detectors and bottom-up excitation and inhibition of word detectors by letter detectors. Readout can occur from either the letter or the word level. In this model the word advantage over single letters under traditional patterned mask conditions is attributed to differential effects of the mask at the letter and word levels. The features in a patterned mask disrupt the letter-level representations but do not replace them with new activations as long as the mask does not contain letters, so that any pattern of activity that has been generated at the word level is allowed to persist longer in the face of masking and has a greater chance of being read out than a pattern of activity at the letter level. Johnston and McClelland's model predicts that if the mask contained letters, the word advantage would be largely eliminated, because the new letter activations caused by the letters in the mask would inhibit active word-detectors.

Johnston and McClelland tested this prediction and found support for it. That is, the presence of letters in the mask strongly disrupted forced-choice performance on letters in words. In contrast, the presence of letters in the mask hurt performance on single-letter displays very little. As a result the word advantage over single letters was reduced

when a mask containing letters was used instead of a nonletter patterned mask.

Our model differs from the Johnston and McClelland (1980) model in that there are top-down and within-level interactions as well as bottom-up interactions. These additional interactions have allowed us to account for the word-superiority effect and a variety of other phenomena without postulating readout from the word level. The word advantage under normal patterned-mask conditions is attributed to feedback from the word level, which strengthens the activations at the letter level. From this we might expect that the presence of letters in the mask would be disruptive to single letters alone and letters in words. In fact with either the standard parameters or the revised parameters needed to account for the pseudoword-enhancement effect, there is little effect of the presence of letters in the mask on either words or single letters. Although the letters in the mask should interfere with response selection, the nodes for the mask letters do not become activated strongly enough to influence response selection until after the peak of the readout function has already passed. However as we noted at the end of Part 1 (McClelland & Rumelhart, 1981), there are other reasons to suppose that there would be readout from the word level at least some of the time. Thus it may be reasonable to admit the possibility that readout may occur from either the letter or the word levels.

Still our model is not completely incapable of handling Johnston and McClelland's (1980) findings even without assuming readout from the word level. Some modifications to the parameters need to be made, however. In order for letters in the mask to make a difference, the visual input must be strong enough to drive the letter detectors to near-ceiling activation values very quickly. This can be done simply by increasing the feature-to-letter excitation and inhibition parameters by a factor of eight. We also need to introduce a ceiling on the maximum amount of inhibition that the feature level can exert on a letter-level node (a value that works well is .55).

Under these conditions, when the target is a single letter, the mask still clears the letter activations very quickly (see Figure

20). However, when the target is a word, the feedback maintains the activations of the letters in the word for a longer period of time, thereby increasing the probability of correct readout. Whether the mask contains letters makes little difference if the activations produced by the target are not being supported by feedback, because in this case the feature-to-letter inhibition drives the letter detectors back down rapidly, causing the new activations produced by the letters in the mask to occur too late to make any difference. However, when there is feedback, the letter activations caused by the target persist long enough for the new activations produced by the mask at the letter level to make a difference. In this case the letters in the mask produce new activations before the output for target letters reaches its maximum strength. These new activations compete with the old ones produced by the target to reduce the probability of correctly encoding the target letter. A second effect of the new letters is to inhibit the activation of the word or words previously activated by the mask. This indirectly results in an increase in the rate of decay of the target letters because their top-down support is weakened. If the mask actually contains a word, it will also eventually produce new activations at the word level. However this effect does not actually come into play until after the peak of the output function has already passed, so it has no effect on performance. In fact Johnston and McClelland (1980) found no difference between masks containing words and masks containing sequences of unrelated letters.

The simulation results shown in Figure 20 were produced using the strong value (.21) of letter-to-word inhibition, in addition to the parameter changes mentioned above. The strong letter-to-word inhibition maximizes the effect of letter masks on words by allowing the letters in the mask to inhibit strongly the word detector that is maintaining the activation of the detectors for letters in the target word.

We do not wish to leave the reader with the feeling that we have been entirely successful in accounting for Johnston and McClelland's findings without requiring readout from the word level. For one thing the simulation shown in the figure has only

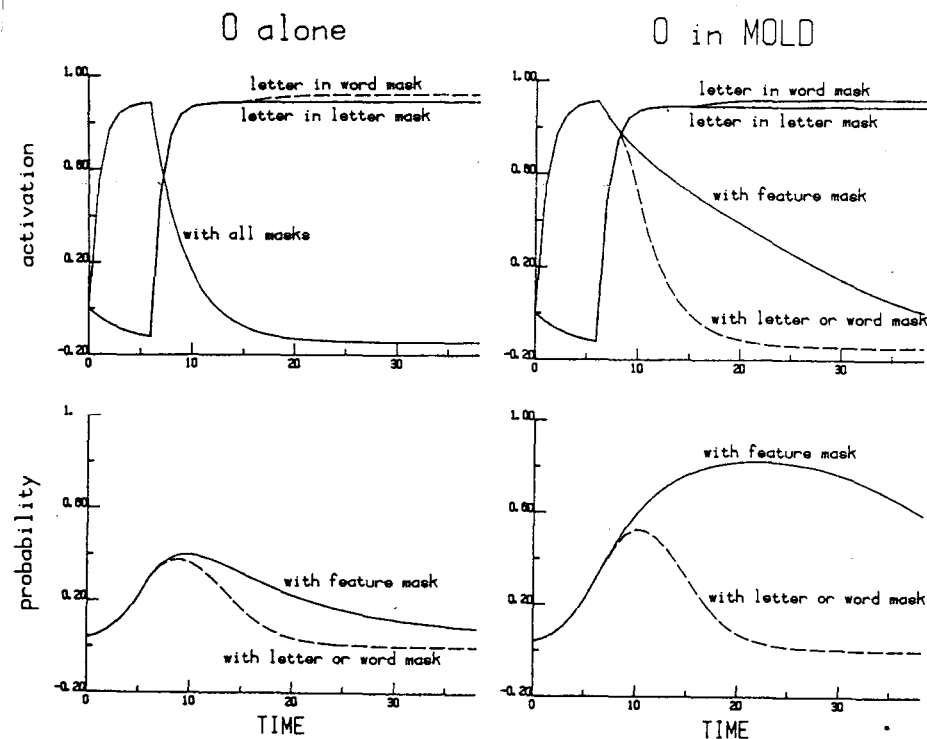


Figure 20. Activation functions (top) and output-probability curves (bottom) for the letter O, both alone (left) and in the word MOLD (right), with feature, letter, and word masks.

been applied to the word MOLD, although all four letters were tested. For another the parameter values we have used to get this simulation to work do not account for the effect of contextual enhancement either with words or pseudowords. Finally the assumption of a maximum value for feature-to-letter inhibition is not independently motivated and we would not wish to defend it. We simply intend this example to indicate that it may eventually prove possible to accommodate Johnston and McClelland's findings in a model like ours without requiring readout from the word level.²

On the Falsifiability of the Model

Although the original parameterization of the model set forth in Part 1 (McClelland

& Rumelhart, 1981) is sufficient to account for several basic findings in the literature, and for the contextual enhancement effect with words, we have found it necessary to modify these parameters to account for the enhancement effect with pseudowords. Further, to account for the findings of Aderman and Smith (1971) and Carr et al. (1978) on the effects of expectations on performance with words and pseudowords, we have been forced to assume that subjects have control

² One reason we have not attempted to account for Johnston and McClelland's (1980) findings more fully is that the conditions under which the interaction of target and mask type can be obtained are not clear at the present. Massaro (Note 3) found no interaction in a similar experiment. The visual display conditions he used were somewhat different from those of Johnston and McClelland, and the experiments also differed in some procedural details.

over at least one parameter of the perceptual system. This gives the model an extra degree of freedom in accounting for the data, of course, though this freedom cannot be exercised unless the conditions of the experiment justify it. Finally, the findings of the Johnston and McClelland (1980) experiments on the effects of letter masks are not compatible with the standard version of the model, but even this does not falsify the general spirit of the model, because there are at least two ways it might be modified to accommodate these findings.

Some flexibility is clearly necessary. After all, one of the characteristics of human perception is its flexibility, and any model that failed to provide for this flexibility would be missing important aspects of the phenomena of perception. A troublesome question arises, though. If we permit such flexibility, do we thereby create an unfalsifiable model? Not necessarily.

There clearly are possible empirical results that would embarrass the model, such as superior performance on unrelated-letter strings compared to words or equal-size enhancement effects for all types of contexts. The fact that the model is capable of accounting for a wide range of different findings under a number of different kinds of conditions with only a very small number of experiment-specific parameters attests to its aptness—certainly there are models that would have failed to follow the ups and downs of the results we have simulated, even if they were allowed complete freedom to adjust parameters between experiments. The only reason why the question arises is that the model has had such great success.

We should also add that the model is not completely unconstrained. Quite the opposite. It contains no detectors for letter clusters, no orthographic rules, and no letter-to-sound translation processes such as other investigators have postulated to account for the research findings on the perceptibility of letters in pseudowords. The fact that the model has been able to do so much with so little attests strongly to the power of the simple computational mechanism embodied in it.

In our view the real issue is whether there is any way of distinguishing our model from

other models. Because in our view our model makes its most provocative contribution to the analysis of the perceptual processing of pseudowords, it is worth a special effort to determine whether it can be distinguished from other possible models of pseudoword perception, including those relying on letter-cluster detectors, systems of abstract orthographic rules, and letter-to-sound translation processes. One way to address this question would be to come up with a general prediction that our model makes that distinguishes it from all or at least some other possible models. In the absence of explicit formulations of other models, this has been somewhat difficult to do. Nevertheless we believe that our model does make one prediction that other approaches to pseudoword perception we know about would not have predicted. We turn now to a test of this prediction.

A Facilitation Effect in All-Consonant Strings

According to our model the pronounceability of a letter string does not determine how accurately the letters in the string will be perceived. All that really matters is how strongly the particular arrangement of letters produces partial word activations that feed back and reinforce activations at the letter level. Thus our model suggests that we might be able to find some unpronounceable nonwords that would produce as much facilitation of perception of the letters in them as comparable pronounceable nonwords would. In this section we report an experiment that demonstrates that such unpronounceable consonant strings do in fact exist.

Experiment 10

The idea of this study was to determine whether there exists a class of unpronounceable and orthographically irregular nonword contexts that nevertheless produce considerable facilitation of the perception of a letter in them.³ Strings of four consonants are clearly orthographically regular and unpro-

³ We are grateful to Mary C. Potter for suggesting this experiment.

nounceable, but in our model they could produce facilitation if the context and the target letter happened to produce partial activations in a number of word nodes. For example, consider the target letter *P* in the string SPCT. This letter occurs in three words that have three letters in common with this display (SPAT, SPIT, and SPOT). The nodes for these words should be activated by the letter string and should produce feedback reinforcing the activation of the *P* node. We would predict, then, that perception of this letter would be facilitated in this context. More generally we would predict that letters that participate in three-letter partial matches with several words should be facilitated, even if the strings consist entirely of consonants. Pronounceability and orthographic regularity per se should make little difference. To test this prediction we tested accuracy of perception of letters in wordlike consonant strings like SPCT and compared performance on these trials to performance on two other types of items: pronounceable pseudowords (e.g., SPET) and nonwordlike consonant strings (e.g., XPQJ).

Method

Stimuli. The stimuli consisted of 20 groups. Each group consisted of a pair of wordlike four-letter consonant strings (like SLCT-SPCT), a pair of pronounceable pseudowords (SLET-SPET) and a pair of nonwordlike consonant strings (SLQJ-SPOJ). The two members of each pair differed from each other by a single letter, and within a group the same two letters differentiated all three pairs of items. Over groups the differing letter occurred in each of the four serial positions equally often. The differing letter between the members of a pair was, of course, the target letter tested in the forced-choice test. Each wordlike consonant string matched at least three words in all but one letter (e.g., SPCT matches the words *spat*, *spit*, *spot*, and *sect* in all but one letter), and the target letter participated in at least two partial matches in every case. As in this example it was possible to make at least one word from each item by replacing the second letter with a vowel, and to make at least one word by replacing the third letter with a vowel. It was never possible to make a word by replacing either end letter with a vowel. For items with the target letters in Positions 1 and 4, this meant that all the words that had three letters in common with the four letters in the string included the target letter and therefore would tend to reinforce its activation. For items with the target letter in Positions 2 and 3, this meant that there was always at least one item that had the three context letters and not the target letter in common with the word shown. Care was taken to ensure that there was only one such word; there were always at least two words that matched

the string shown in the target letter position and two other positions.

From each pair of wordlike consonant strings, we generated a pair of pronounceable pseudowords and a pair of nonwordlike consonant strings as follows: The pseudowords were constructed by replacing one of the context letters in each pair with a vowel. The replaced letter was always one of the two internal letters (e.g., to go with the pair SLCT-SPCT, the pseudoword pair was SLET-SPET), and the resulting string was never an actual word in English. For target letters in the first and second position, the third letter was always replaced. For target letters in the third and fourth position, we wanted to replace the second letter with a vowel but were unable to avoid making words in three cases with fourth-position target letters and so had to replace the third letter with a vowel in these cases. For similar reasons it was necessary to change the final consonant as well as the third letter in the pair SLRT-SPRT to make SLAD-SPAD.

The matched nonwordlike consonant strings were constructed by replacing the three context letters in each wordlike pair with a permutation of the letters *Q*, *X*, and *J*. These letters were chosen to minimize the number of words that would match the four-letter strings in three or even two letter-positions.⁴

In addition to these test materials, practice and filler stimulus-pairs were selected from the pronounceable pseudowords used in Experiment 6.

Procedure. The procedure followed was identical to that used in Experiments 4 and 7-9, with the following changes: The stimuli were only presented in the normal presentation condition—that is, all four letters were turned on and off together, followed immediately by the \square mask. Stimuli were arranged into 24 blocks of 16 trials. There was no break between blocks, but the program automatically recalculated the optimal exposure duration to achieve an overall performance level of 75% correct after each block of trials. The first four blocks were practice trials consisting only of filler pseudoword-items. Each of the remaining blocks contained in random order 10 pseudoword fillers, 2 of the wordlike consonant strings, 2 of the matched pseudowords, and 2 of the nonwordlike consonant strings. Each serial position was tested equally often in each block, and over adjacent pairs of blocks each serial position was tested equally often in each type of material. In the first 10 blocks of experimental trials, one member of each pair of experimental items was presented. The other member of each pair was presented in the second 10 blocks of experimental trials. There was no repetition of filler materials.

A different randomly determined stimulus list was constructed for each subject. Stimulus lists were constructed in yoked pairs so that items that appeared in the first 10 blocks for one subject appeared in the second 10 blocks for the other and vice versa.

Results

The wordlike consonant strings produced 16% more accurate performance than the

⁴ The complete list of stimuli used in this experiment is available from either author.

nonwordlike consonant strings, and there was virtually no difference in overall accuracy between the wordlike consonant strings and the pseudoword strings (see Table 6). Of course we cannot conclude that letters in our wordlike consonant strings are actually just as easily perceived as comparable letters in pseudowords, because the 95% confidence interval around the difference between the two conditions is $\pm 4.6\%$. But it is clear that both the pronounceable pseudowords and the wordlike consonant strings produce a highly reliable advantage over letters in nonwordlike consonant strings ($p < .001$ for both comparisons).

The performance of the model on the stimuli used in this experiment parallels the actual results. The simulation was run twice, once with the standard set of parameters given in Part 1 (McClelland & Rumelhart, 1981) and once with the revised set used to account for the pseudoword-enhancement effect. The mask was presented after 15 cycles, with readout after Cycle 16 for all three material types. Using the standard parameters the model produced a slightly smaller advantage for the wordlike consonant strings over the nonwordlike strings than was actually observed, but as in the actual data there was virtually no difference between the wordlike consonant strings and the pronounceable nonwords. With the revised parameters the results were nearly identical. These parameters produced a 2% advantage for the pronounceable pseudowords over the wordlike consonant strings, and a 9% advantage for the latter over the nonwordlike consonant strings.

The serial-position curves produced in the experiment are illustrated in Figure 21. These curves seem at first glance to suggest that there is, in fact, a processing difference

between pronounceable pseudowords and all-consonant strings. The serial-position curve is nearly flat for pronounceable pseudowords but falls off dramatically from the first to the second position for the other two types of material. The overall equality of the wordlike consonant strings and the pronounceable pseudowords arises from the fact that the consonant strings start higher in the first position before the steep drop in performance for the second-position items. Taken at face value these results suggest that pronounceable pseudowords, but not consonant strings, may be perceived as wholes. However it turns out that the model can account for the general properties of these serial-position curves assuming the same readout processes for all conditions. The serial-position curves generated by the model made use of the two additional assumptions introduced previously in accounting for the serial-position results of the enhancement experiments. First, the rate of processing of the different letters in the display varied over letter position. The rate of processing the first letter was set to 1.6 times the normal rate, and the rates for the other letters were set to 1.05. Second, it was assumed that readout occurred in an outside-in order. The readout for the first and last letters was coincident with the offset of the mask, whereas the readout for the second and third letters occurred four cycles later. The mask replaced the target display on Cycle 14.

The results were nearly identical for both the standard and revised parameters. The only difference was that the revised parameters produced a 2% pseudoword advantage over wordlike consonant strings whereas the standard parameters produced no difference. The curves shown in the figure are for the standard parameters.

Why does the model produce different serial-position curves for the different types of materials? Part of the answer lies in the distribution of partially matching words for the pronounceable pseudowords and the wordlike consonant strings. As illustrated in Table 7, the target letters in the pseudoword items have far more friends, particularly for the second and third letters, than the wordlike consonant strings. At the same time the consonant strings have fewer enemies in the end letter positions.

Table 6
Probability of Correct Forced Choice for Obtained and Simulated Results From Experiment 10

Context	Obtained	Simulated
Wordlike consonant strings	.78	.79
Pronounceable nonwords	.78	.79
QXJ context	.62	.68

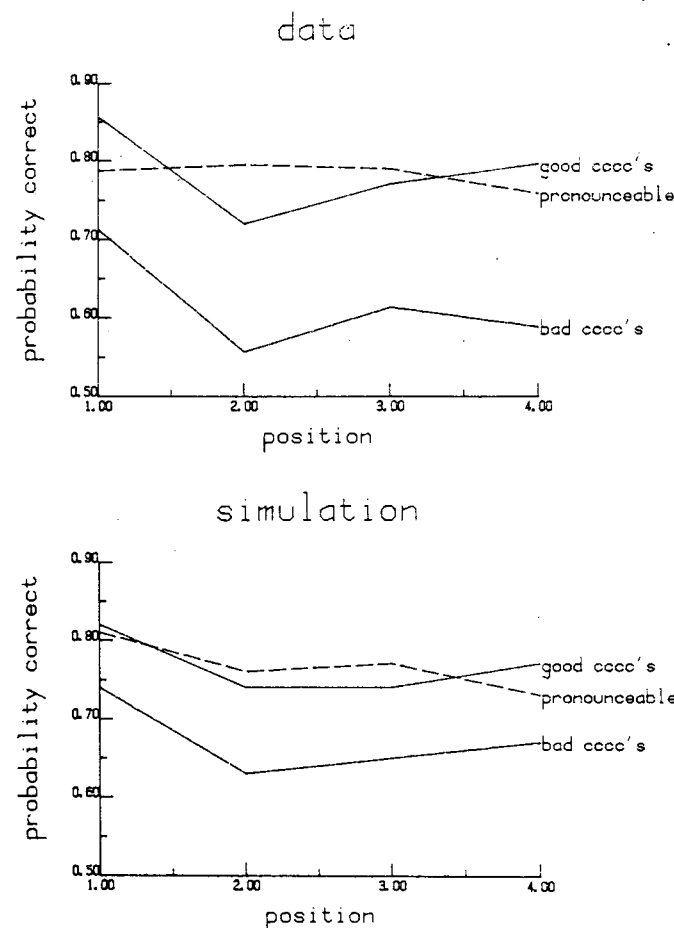


Figure 21. Serial-position curves for the wordlike consonant strings, pronounceable pseudowords, and nonwordlike consonant strings from Experiment 10. (In the simulation the first letter was given a stronger weight than the other letters, and readout occurred earlier for end letters than for internal letters.)

Discussion

Experiment 10 confirms the predictions of the interactive activation model. Our simulations predicted that performance would be approximately the same for the pronounceable pseudowords and the wordlike consonant strings, and that both would produce more accurate performance than the nonwordlike consonant strings. The results came out exactly as predicted.

Several alternative interpretations of the pseudoword advantage over unrelated-letter strings appear to be inconsistent with these findings. Clearly models in which perceptibility depends on the pronounceability of pseudowords or on the fact that they are orthographically regular (in the sense that they are candidates to be words in English if a new one is needed) would not predict this pattern of results. Our wordlike consonant strings are neither pronounceable nor

Table 7
Friends and Enemies of the Critical Letter for the Stimuli Used in Experiment 10

Stimuli	Position			
	1	2	3	4
Wordlike Consonant Strings				
Friends	4.400	2.700	3.300	4.800
Enemies	.000	1.000	1.000	.000
Pronounceable Pseudowords				
Friends	5.400	5.400	8.200	5.400
Enemies	1.000	.600	.800	1.200

orthographically regular, yet they produce an advantage just as clearly as pronounceable, orthographically regular stimuli do.

Of course it would be reasonable to suggest that orthographic regularity and pronounceability are matters of degree, and to point out that our wordlike consonant strings are considerably more pronounceable and orthographically regular than our nonwordlike consonant strings. From these observations models that stress orthographic regularity or pronounceability may easily be formulated that will predict an advantage for our wordlike consonant strings over nonwordlike consonant strings. Indeed differences in perceptibility among strings that are not strictly pronounceable have been obtained before (Spoehr & Smith, 1975), as predicted from a model that accounted for the pseudoword advantage in terms of the construction of a phonological code. However the Spoehr and Smith model also predicted that completely pronounceable items like our pronounceable pseudowords would have an advantage over even the most pronounceable consonant strings. Though they found evidence of this, we found no such effect in our experiment. We cannot, of course, be sure of the reason for the discrepancy. However our materials were selected to maximize the benefit the target letter might receive from partial activations of words, whereas theirs were not.

More generally, any account of the pseudoword advantage that predicts that perceptibility correlates with orthographic regularity or pronounceability (including the

interpretation offered by McClelland & Johnston, 1977) would lead us to expect some advantage of our pseudowords over our wordlike consonant strings, and no such advantage was obtained. Although proponents of such views could argue that our failure to detect such an effect was due to some insensitivity of our experiment, our data leave no reason to prefer such models over the interactive activation model, in which the advantage for pronounceable pseudowords over unpronounceable nonwords is due not to pronounceability or orthographic regularity but to feedback generated from the partial activation of representations of words.

An alternative to the notion that orthographic regularity or pronounceability determines perceptibility is the view that perceptibility varies with the frequency of occurrence of multiletter substrings such as bigrams and trigrams. Such a view is certainly consistent with our finding that letters in our wordlike strings are perceived more accurately than letters in our nonwordlike consonant strings. However the bigram frequencies are certainly higher (especially across Positions 2 and 3) in our pronounceable strings than in our wordlike consonant strings, and the trigrams in the wordlike consonant strings are almost all completely unfamiliar. Thus it is difficult to see how a model explaining the pseudoword advantage on the basis of substring detectors would not end up predicting superior performance for the pronounceable strings than the wordlike consonant strings would.

A final alternative is the notion that perceptibility is based only on positional frequencies of occurrence of single letters. A variety of investigators have reported positional frequency effects in perceptual-accuracy studies and related tasks (Mason, 1975; Massaro, Venezky, & Taylor, 1979; McClelland, 1976; McClelland & Johnston, 1977). However this factor has rarely been thought to be solely responsible for perceptual differences between words and pronounceable pseudowords, and Massaro et al. (1979) demonstrated that conformity to letter co-occurrence rules was correlated both with judgments of "similarity to real words in English" and with performance in a perceptual accuracy task similar to the Reicher

(1969) task, even after positional frequency effects had been taken into account. Our model is, of course, quite consistent with some correlation of positional frequency and accuracy because positional frequency in words is strongly related to the number of words a letter might help to activate. A model that attributed positional frequency effects directly to greater sensitivity of position-specific letter detectors for frequent letters in that position would not (at least without further assumptions) explain the fact that positional frequency of context letters increases accuracy of performance on the target (Johnston, 1978; McClelland & Johnston, 1977). In contrast, we would definitely expect such effects in our model.

Of course we are not claiming that there can be no model other than ours that is consistent with the results of Experiment 10 and all of the other experiments we have examined in these two articles. It does seem, however, that the present experiment lends considerable support to the general approach we have taken in accounting for the perceptual advantage of words and pseudowords and provides little comfort to alternative approaches.

It may be suggested that we can test our model by looking for differences between orthographically regular pseudowords that depend on the number of four-letter-word friends and enemies of the particular item in question. However a failure to find such differences would be inconclusive. Though the specific version of the model we have simulated might be ruled out in this way, it would not only be possible, but also highly sensible to argue that we had simply failed to include all of the relevant word-level knowledge in our model. After all it is likely that three-, five-, and even six-letter words might become partially activated when four-letter displays are shown. Such words would probably fill in the gaps in the coverage of the four-letter words and thereby permit letters with few friends among the four-letter pseudowords to gain additional support, thereby weakening the experiment considerably.

It should be noted in view of these remarks that the prediction of the model that we have just tested would not be much altered by

partial activations of words of other lengths, because both the pronounceable pseudowords and the partially matching consonant strings already have large numbers of friends, whereas their nonwordlike consonant strings are not similar to any words of any length.

Some Extensions of the Model

Relating Retinal Position to Position in a Word

Until this point we have completely ignored the fundamental problem of how visual features that are initially registered by receptors in particular locations on the retina are mapped onto the four-letter slots in our model. This is not a trivial matter because it is possible for us to read words in print of various sizes, in any position on the retina, provided only that the resolution is sufficient. Hinton (Note 4) has recently proposed a general scheme that employs interactive processing of the general sort we have outlined here to carry out the whole range of transformations that might be involved in such a mapping. These include rotation (in three dimensions), translation (also in three dimensions), and size adjustment. The basic idea is that each possible mapping is associated with a unit that modulates the extent to which a particular feature of the retinal display activates a particular unit in the canonical activation network for perception of patterns. In our case, for example, these mapping units would determine to what extent features in a particular location in the retinal array would activate units for features in particular positions relative to the beginnings and endings of words. Initially each possible mapping is open, but as processing continues the mappings that are most consistent with some stable higher order perceptual structure are strengthened and come to dominate all of the other mappings, thereby effectively closing all the activation paths associated with all the other mappings. One implication of this notion is that information about position and information about the identity of letters may become separated in the perceptual system if the set of retinal features for a particular letter end up being mapped onto the right set of canonical fea-

tures but in the wrong canonical position. In fact experiments using the full-report or probed-report procedure show that subjects often rearrange letters in their reports, indicating that they have picked up the identity of the letters shown without necessarily picking up their order (Estes, 1975; McClelland, 1976).

We have not attempted to incorporate Hintikka's ideas fully into our model. However it is worth considering the possibility that information presented in one location might activate detectors in a range of locations, rather than just simply in one fixed position. Perhaps there is a region of uncertainty associated with each feature and with each letter. If so a given feature in a given input position would tend to activate units for that feature in positions surrounding the actual appropriate position. As a result partial activations of letters from nearby positions would arise in a particular position along with the activation for the letter actually presented. Similarly this same input might partially activate word units for words with that letter in neighboring positions. In a scheme such as this, the role of feedback from higher levels would not only be to reinforce letters consistent with some known pattern or combination of known patterns, but also to reinforce the activations of these letters in particular positions at the expense of other positions. It should be clear, then, how this scheme could cause transposition errors, especially in those cases where the transposition makes a more wordlike string than the original does (e.g., TAED → TEAD). Thus it appears that a scheme of this sort offers a plausible account for the finding that irregular nonword strings are often reported with letters transposed if the transposition will produce regular strings (Estes, 1975; c.f., experiment by Stevens reported in Rumelhart, 1977).

A mechanism of this type would have the property of "smearing" the pattern of activation produced by a stimulus of one length over a somewhat longer array of possible positions. One side effect of such smearing would be that it would tend to produce activations of words of other lengths besides those corresponding to the actual length of the input. Such activations could, of course,

generate feedback, supporting the letter-level activations that got them going. Such support could be very valuable, especially to the perception of pseudowords. Such support would make performance on pseudowords less dependent on the details of the set of four-letter words and, in that sense, more robust.

Extensions to Other Domains

In the previous sections of this paper, we have focused on producing accounts of the perception of letters in the context of words and pseudowords. However the modeling framework we have been working with is potentially much more general than this. At the outset we proposed a more general framework for the processing of words and pseudowords in either the visual or the auditory modality. Figure 22 (Figure 1 from Part 1; McClelland & Rumelhart, 1981) shows the general view with which we began. Here we have, in addition to the visual processing system, a speech processing system, including an acoustic-feature level and a

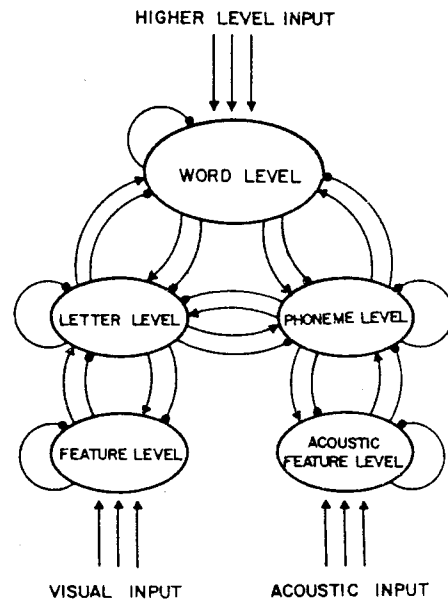


Figure 22. Full version of the interactive activation system for visual and auditory word recognition.

phonological level. The figure also provides for input from higher contextual levels. In this section we discuss a number of results that we believe could be accounted for with an extended version of our model, incorporating the processes represented in Figure 22.

Use of Context in Word Recognition

It is, of course, a well-known fact that context preceding the visual presentation of a word can influence identification of the word (Tulving & Gold, 1963; Tulving, Mandler, & Bauml, 1964). The results of these experiments have been accounted for by the logogen model of McClelland (1977). Our model is quite like the logogen model in various ways, as we pointed out in Part 1, and it is clear that it would account for the facilitation and interference effects of appropriate and inappropriate context. As in the logogen model, we would simply imagine that a context would tend to prime the nodes for words consistent with it. Such words would tend to benefit from this priming. The interference with performance of words inconsistent with the context could be accounted for in various ways. One possibility is that the contextual inputs directly inhibit the nodes for words not consistent with the context. However direct inhibition may not be necessary to account for the effect, because our model provides two other possible sources of interference. Priming one set of words would result in relative inhibition of nodes for other words, owing to the lateral interference mechanism, provided the context was strong enough to drive the activations of nodes for contextually appropriate words above the interaction threshold. Furthermore the process of selecting a response takes into account the strengths of all of the logogens, and active logogens other than the correct one have the effect of reducing the probability that the correct response will be chosen.

Because we have considered performance in Reicher's (1969) forced-choice paradigm so extensively, it seems appropriate to comment on the effect of linguistic context on performance in this paradigm. Though there are no published studies on these effects,

there are two unpublished findings. One of these comes from a study by Hale and Johnston (Note 5). They compared forced-choice performance for letters in words when the presentation was preceded by a context sentence similar to those used by Tulving and Gold (1963). On half the trials the item shown was appropriate to the context; on the other half, the item was an inappropriate word differing by a single letter from a word that would fit. For example, one of the sentences might have been:

I like to follow and I hate to ____

For this context the word pair might have been LEAD-DEAD. The subject's task was to choose between the two alternatives. The results were that the context had a large effect, biasing performance in the direction of the contextually appropriate word, but it did not have a noticeable effect on the average probability correct (or on d'), compared to a condition in which the same items were presented with no preceding context. The other piece of data comes from an unpublished experiment by one of us (Rumelhart). In this study stimulus triples consisting of a prime word and two target words (both semantically related to the prime and differing from each other by a single letter) were used (e.g., WAR-ARMS-ARMY). As in the Johnston (1978) experiment, the prime was followed by a brief, masked presentation of one of the two alternatives, and this was followed by a forced choice between the differing letters. Also as in the Johnston experiment, there was hardly any overall effect of the prime on accuracy, compared to a control condition with an inappropriate prime. This time accuracy improved a little bit when an appropriate prime was used, but the difference was only 2% in the forced choice, and it was not significant.

We have run informal simulations of these two experiments. To simulate the effects of the prime, we simply imposed a constant input to one or more word nodes as appropriate and left it on for a few cycles until the activation of the node stabilized. Then we presented the target stimulus to the model as before. We found that when the prime preactivated only one word (e.g., LEAD, as in the example of a possible sentence from the

Johnston experiment), there was a considerable bias toward choosing the forced-choice alternative letter consistent with the primed word. This bias helped performance on those trials when LEAD was actually shown to the model, but it had an equal effect in the opposite direction when LEND was shown. When the prime preactivated both choice alternatives, there was only a very slight effect—as in the actual data, there was a small benefit, which did not grow larger than about 2% even with a very substantial prime.

Pronouncing Words and Pseudowords

Glushko (1979) has recently presented several important findings that suggest that a model structured very much like our model may underlie the process of constructing pronunciations for both words and pseudowords. The central finding of Glushko's studies is that the time it takes to pronounce an orthographically regular probe nonword (such as MAVE) depends on the pronunciations of words that have the same last three letters as the probe. Consider MAPE in contrast to MAVE. Of the English monosyllables ending in APE, all of them rhyme with TAPE. Such pseudowords are said to have consistent neighborhoods. On the other hand, of the English monosyllables ending in AVE, one (HAVE) has a pronunciation that is different from that of all of the others (which rhyme with CAVE). Glushko found that pseudowords (and, in fact, words) with inconsistent neighborhoods are pronounced more slowly than corresponding items with consistent neighborhoods.

To account for this result, Glushko proposed an alternative to the usual notion that pseudowords are pronounced by applying explicit pronunciation rules. Instead he suggested that pronunciation proceeded by partial activation of the pronunciations of all the words in the neighborhood of the probe pseudoword, followed by synthesis of a pronunciation from these partial activations. Such a model is, of course, consistent with the spirit of our model and, in fact, as mentioned previously, was part of the inspiration for our interpretation of the perceptual advantage for letters in pseudowords. Simulation of this and other results obtained by

Glushko, however, await the construction of a version of our model that includes stored information about the pronunciations, as well as the spellings, of words in English.

Perception of Speech

The processing structure we have explored in this paper may have general utility in modeling other psychological phenomena beyond those concerned with the perception and pronunciation of visually presented words. One very promising extension of the model would be into the area of speech perception.

The role of context in speech perception is, of course, very well established (see Foss & Blank, 1980, for a recent review). The most striking example is the phenomenon of phonemic restoration. When listening to words in context, subjects perceive whole phonemes that have been replaced by a noise, a cough, or a tone as if the replaced version had actually been spoken clearly and normally (Warren, 1970; Warren & Obusek, 1971). A recent series of studies by Samuel (1979) has demonstrated that restorations are more likely when words occur in semantically predictive contexts, when the phoneme to be restored occurs later in the word, and when the phoneme to be restored occurs in a word rather than in a pronounceable nonword. Related findings in slightly different tasks have been reported by Cole (1973) and Marslen-Wilson and Welsh (1978). That the phonemic-restoration effect and related phenomena are perceptual rather than merely a matter of postperceptual biases is indicated by the fact that restorations occur phenomenologically even when subjects are prewarned that speech sounds will be excised. Indeed the Samuel (1979) experiments indicate that context can actually reduce subjects' ability to distinguish between a complete word with an extraneous sound superimposed on one of the phonemes and a complete word with an extraneous sound replacing one of the phonemes.

Our general modeling-framework provides a natural way of accounting for these context effects in speech perception. We might imagine, for example, that speech per-

ception takes place within a multilevel system like the one shown in Figure 22. Assuming that readout of the results of perceptual processing occurred at the phonological level, we would explain the results reviewed above by supposing that this readout process was guided both by the acoustic features of the input itself and by top-down activation from higher levels through the word level to the phonological level.

Although the speech perception model would be quite similar to the model we have described for the perception of printed words, there would be several differences as well. One of these would be that the information in the input supporting many of the phonetic features would be spread over the nominal locations of several of the phonemic segments in the input (Studdert-Kennedy, 1976). Another important difference would be that the speech signal unfolds over time, so that the information from the initial portions of a word would be available for processing before information from later portions of the same word would be. For a word presented without any prior context, this means that the beginning of a word arrives in an unprimed system, whereas the later portions of the word arrive in the system after activations at the phonological and lexical levels have had a chance to become established and lexical activations have had a chance to begin activating phonological segments for later portions of the word. Whether there would be other fundamental differences between the model for speech perception and the model for visual perception of words remains to be seen.

Conclusion

The model we have explored in this paper attempts to explain the role of familiar context in perception in terms of simple excitatory and inhibitory interactions among large populations of very simple neuronlike units. In these respects the model is a part of a recent trend toward trying to apply neural or neurallike models to cognitive processes (Anderson, 1977; Grossberg, 1980; Hinton, 1977; see Hinton & Anderson, 1981, for a recent review).

We have found our simulation method to

be exceptionally useful for the study of processing systems of the kind we have described here. Time and again during the development of this model, we found that our intuitions about how the model would behave were incorrect. The use of such simulations may be the only way to get a sufficient handle on complex interactive process such as these to be able to make any unequivocal claims about the behavior of the system in a particular situation.

We hope that our explorations in this new domain will contribute to the growing feeling that it is a fertile one. The model we have constructed in this framework appears to provide a very close account of many of the major phenomena in word perception, including some new findings that we have presented on the way contextual inputs influence perceptual processing. The model appears also to provide a plausible framework for accounts of the perception of visually presented words in linguistic context, for the perception of phonemes in speech, and for the translation of written words and pronounceable nonwords into a phonological code.

We have focused our analysis on the visual processing of words, though we have tried to indicate that the framework is much broader than this. We did not choose to focus on word perception because we believe that this task requires a unique mode of processing. Rather we focused on these phenomena as especially well-studied examples of processes that are ubiquitous in the human information-processing system. There is a wealth of detailed experimental observations that have served to constrain and inform our model-building enterprise. In addition to the extensions considered above, we are already at work constructing similar models for motor production and for concept abstraction.

Perhaps the single most unique feature of our account is that we have offered a single mechanism for the processing of both familiar stimuli and items that are structurally similar to familiar stimuli but that are not themselves familiar wholes. Specifically the mechanism has been used to account for the perception and pronunciation of both familiar words and novel pseudowords. Most previous models that have attempted to account

for perception of novel but structurally regular stimuli have relied on the use of a stored system of rules. We have shown how, through the use of interactive processes, the mere activation of stored representations of familiar patterns can suffice, at least to account for the perception of letters in novel pseudowords. There are fundamental problems to be overcome before such a mechanism can be applied to several other instances of processing novel, structurally regular patterns, and it is not now clear whether it will be necessary in some cases to postulate the use of stored systems of abstracted rules. However our explorations suggest that it may be fruitful to continue exploring the possibility that other types of apparently rule-governed behavior may be accounted for by synthesis of stored knowledge about individual cases.

Reference Notes

1. Johnston, J. C. Personal communication (undated).
2. Rumelhart, D. E., & McClelland, J. L. *An interactive activation model of the effect of context in perception, Part 2* (Chip Technical Report #95). San Diego, Calif.: Center for Human Information Processing, University of California, San Diego, 1980.
3. Massaro, D. W. Personal communication (undated).
4. Hinton, G. E. *A parallel computation that assigns canonical object-based frames of reference*. Cambridge, England: MRC Applied Psychology Unit, 15 Chaucer Road, Cambridge CB2 2EF, England, 1981.
5. Hale, B., & Johnston, J. C. Personal communication, July 1981.

References

- Aderman, D., & Smith, E. E. Expectancy as a determinant of functional units in perceptual recognition. *Cognitive Psychology*, 1971, 2, 117-129.
- Anderson, J. A. Neural models with cognitive implications. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension*. Hillsdale, N.J.: Erlbaum, 1977.
- 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.
- Cole, R. A. Listening for mispronunciations: A measure of what we hear during speech. *Perception & Psychophysics*, 1973, 13, 153-156.
- Eriksen, C. W., & Rohrbaugh, J. Visual masking in multielement displays. *Journal of Experimental Psychology*, 1970, 83, 147-154.
- Estes, W. K. Memory, perception, and decision in letter identification. In R. L. Solso (Ed.), *Information pro-*

- cessing and cognition: The Loyola Symposium*. Potomac, Md.: Erlbaum, 1975.
- Estes, W. K., Allmeyer, D. H., & Reder, S. M. Serial position functions of letter identification at brief and extended exposure durations. *Perception & Psychophysics*, 1976, 19, 1-15.
- Foss, D. J., & Blank, M. A. Identifying the speech codes. *Cognitive Psychology*, 1980, 12, 1-31.
- Glushko, R. *The psychology of phonography: Reading aloud by orthographic activation and phonological synthesis*. Unpublished doctoral dissertation, University of California, San Diego, 1979.
- Grossberg, S. How does the brain build a cognitive code? *Psychological Review*, 1980, 87, 1-51.
- Hinton, G. E. Relaxation and its role in vision. Unpublished doctoral dissertation, University of Edinburgh, Scotland, 1977.
- Hinton, G. E., & Anderson, J. A. (Eds.). *Parallel models of associative memory*. Hillsdale, N.J.: Erlbaum, 1981.
- 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. Experimental tests of a hierarchical model of word identification. *Journal of Verbal Learning and Verbal Behavior*, 1980, 19, 503-524.
- Kucera, H., & Francis, W. *Computational analysis of present-day American English*. Providence, R.I.: Brown University Press, 1967.
- Marslen-Wilson, W. D., & Welsh, A. Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, 1978, 10, 29-63.
- Mason, M. Reading ability and letter search time: Effects of orthographic structure defined by single-letter positional frequency. *Journal of Experimental Psychology: General*, 1975, 104, 146-166.
- 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. L. On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, 1979, 86, 287-330.
- 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 context effects in letter perception: Part 1. An account of basic finding. *Psychological Review*, 1981, 88, 375-407.
- Miller, G. A., Bruner, J. S., & Postman, L. Familiarity of letter sequences and tachistoscopic identification. *Journal of Genetic Psychology*, 1954, 50, 129-139.
- Morton, J. Interaction of information in word recognition. *Psychological Review*, 1969, 76, 165-178.
- Reicher, G. M. Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 1969, 81, 274-280.

- Rumelhart, D. E. Toward an interactive model of reading. In S. Dornic (Ed.), *Attention and performance VI*. Hillsdale, N.J.: Erlbaum, 1977.
- Rumelhart, D. E., & Siple, P. The process of recognizing tachistoscopically presented words. *Psychological Review*, 1974, 81, 99-118.
- Samuel, A. G. Speech is specialized, not special. Unpublished doctoral dissertation, University of California, San Diego, 1979.
- 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.
- Studdert-Kennedy, M. Speech perception. In N. J. Lass (Ed.), *Contemporary issues in experimental phonetics*. New York: Academic Press, 1976, 243-293.

- Tulving, E., & Gold, C. Stimulus information and contextual information as determinants of tachistoscopic recognition of words. *Journal of Experimental Psychology*, 1963, 66, 319-327.
- Tulving, E., Mandler, G., & Bauml, R. Interaction of two sources of information in tachistoscopic word recognition. *Canadian Journal of Psychology*, 1964, 18, 62-71.
- Warren, R. M. Perceptual restoration of missing speech sounds. *Science*, 1970, 167, 393-395.
- Warren, R. M., & Obusek, D. J. Speech perception and phonemic restorations. *Perception & Psychophysics*, 1971, 9, 358-362.

Received March 11, 1981

Revision received July 31, 1981 ■