

2

DTIC FILE COPY

AD-A219 113

CONTEXT EFFECTS IN LETTER PERCEPTION:
A COMPARISON OF TWO THEORIES

Technical Report AIP - 45

Howard B. Richman & Herbert A. Simon

Department of Psychology
Carnegie Mellon University
Pittsburgh, Pa. 15260

The Artificial Intelligence
and Psychology Project

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTE
MAR 14 1990
S E D

Approved for public release; distribution unlimited.

90 03 12 015

CONTEXT EFFECTS IN LETTER PERCEPTION:
A COMPARISON OF TWO THEORIES

Technical Report AIP - 45

Howard B. Richman & Herbert A. Simon

Department of Psychology
Carnegie Mellon University
Pittsburgh, Pa. 15260

30 June 1988

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

DTIC
ELECTE
MAR 14 1990
S E D

This research was supported by the Computer Sciences Division, Office of Naval Research and DARPA under Contract Number N00014-86-K-0678. Reproduction in whole or in part is permitted for purposes of the United States Government. Approved for public release; distribution unlimited.

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION Unclassified		1b RESTRICTIVE MARKINGS	
2a SECURITY CLASSIFICATION AUTHORITY		3 DISTRIBUTION AVAILABILITY OF REPORT Approved for public release; Distribution unlimited	
2b DECLASSIFICATION/DOWNGRADING SCHEDULE			
4 PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 45		5 MONITORING ORGANIZATION REPORT NUMBER(S)	
6a NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b OFFICE SYMBOL (if applicable)	7a NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research	
6c ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213		7b ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000	
8a NAME OF FUNDING/SPONSORING ORGANIZATION Same as Monitoring Organization	8b OFFICE SYMBOL (if applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678	
8c ADDRESS (City, State, and ZIP Code)		10 SOURCE OF FUNDING NUMBERS p4000ub201/7-4-86	
		PROGRAM ELEMENT NO N/A	PROJECT NO N/A
		TASK NO N/A	WORK UNIT ACCESSION NO N/A
11 TITLE (Include Security Classification) CONTEXT EFFECTS IN LETTER PERCEPTION: A COMPARISON OF TWO THEORIES			
12 PERSONAL AUTHOR(S) Richman, Howard B. & Simon, Herbert A.			
13a TYPE OF REPORT Technical	13b TIME COVERED FROM 86Sept15 TO 91Sept14	14 DATE OF REPORT (Year, Month, Day) 1988 June 30	15 PAGE COUNT 35
16 SUPPLEMENTARY NOTATION			
17 COSATI CODES		18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19 ABSTRACT (Continue on reverse if necessary and identify by block number) The purpose of this study is to test whether EPAM (Elementary Perceiver and Memorizer) can explain context effects in letter recognition. EPAM, a model of learning and recognition in the form of a computer program, has successfully explained many aspects of learning and perception in a range of task environments. In 1984, Barsalou and Bower claimed that EPAM could not explain the phenomena in the tachistoscopic perception experiments successfully simulated by the Interactive Activation Model (IAM) of word perception (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982). In this study, using EPAM IV, a revision of the most recent version of EPAM, we show that the human data modeled by the IAM are at least as accurately simulated by EPAM. Close examination of the performance of the two programs shows that the fact that one (EPAM) processes perceptions serially, while the other (IAM) processes them in parallel, plays no essential role in producing the context effects that are observed. The main effects are produced, in both programs, by a feedback of information from word recognition to the process of recognizing letters.			
20 DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS		21 ABSTRACT SECURITY CLASSIFICATION	
22a NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Mayrowitz		22b TELEPHONE (Include Area Code) (202) 696-4302	22c OFFICE SYMBOL N00014

Context Effects In Letter Perception: A Comparison of Two Theories

Howard B. Richman and Herbert A. Simon
Carnegie-Mellon University

Abstract

The purpose of this study is to test whether EPAM (Elementary Perceiver and Memorizer) can explain context effects in letter recognition. EPAM, a model of learning and recognition in the form of a computer program, has successfully explained many aspects of learning and perception in a range of task environments. In 1964, Barsalou and Bower claimed that EPAM could not explain the phenomena in the tachistoscopic perception experiments successfully simulated by the Interactive Activation Model (IAM) of word perception (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982). In this study, using EPAM IV, a revision of the most recent version of EPAM, we show that the human data modeled by the IAM are at least as accurately simulated by EPAM.

Close examination of the performance of the two programs shows that the fact that one (EPAM) processes perceptions serially, while the other (IAM) processes them in parallel, plays no essential role in producing the context effects that are observed. The main effects are produced, in both programs, by a feedback of information from word recognition to the process of recognizing letters.

Progress in cognitive psychology in recent years has produced a number of rather distinct theories to account for the basic processes of perception and cognition. Before we examine the differences between two of these theories, which is the main task of this paper, we need to emphasize what they all have in common:

1. They characterize perceptual and cognitive processes as information processes.
2. They are computational: they are described with such specificity that they can be, and are, programmed for digital computers. Hence, their predictions can be compared with data at a relatively fine level of concrete detail.
3. They are quite general. That is, they model, and make predictions for, a wide range of perceptual and cognitive performances, and not just performance on one task or a few closely similar tasks.

Among the theories that share these characteristics are the connectionist models put forward by McClelland and Rumelhart (1981, 1982), the EPAM theory of Feigenbaum and Simon (1963, 1984), John Anderson's ACT* (1983), and Allen Newell's SOAR (198X). The first of these describes thinking without the use of symbols, as based upon patterns of activation in network-like structures. The other three accept the physical symbol system hypothesis (Newell and Simon, 1976), that intelligence is produced by the processing of symbols in a physical system capable of such processing.

Although all four theories have considerable generality, their ranges are not quite coterminous. SOAR makes perhaps the broadest claims to generality, but it has been elaborated mainly in application to cognitive tasks, and often quite complex ones. In ACT*, the emphasis is

focused upon the structure and operation of semantic memory, and particularly on the role of spreading activation in performance and learning. ACT* does not provide a detailed account of perception.

EPAM focuses primarily on perception and recognition and on learning, but it does not contain a model of the initial, feature-detection phases of perception. Short-term memory plays a central role in EPAM; its model of long-term memory is somewhat incomplete. Hence, it has been most useful in accounting for relatively simple perceptual and learning tasks rather than complex problem-solving tasks. A "package" of EPAM with the compatible models, UNDERSTAND (Simon and Hayes, 1974), and GPS (Newell and Simon, 1972) would comprise a more nearly comprehensive theory of perception and cognition.

The connectionist theory of Rumelhart and McClelland starts with perception (though excluding the initial steps of feature detection), but in other research has been extended to some fairly complex cognitive tasks.

Testing Theories

The existence of these theories and their claims to explain phenomena in domains that overlap to a very considerable extent create a task for contemporary cognitive science. The task is to develop methods for comparing the relative adequacies of the theories for explaining experimental data in situations to which more than one of them applies. This may seem a relatively straightforward affair: simply see which theory fits the data best. Unfortunately, matters are not nearly so simple.

First problem: Each of the theories takes the form of a fairly complex computer program. Various parameters of the program can be adjusted, and the numbers and kinds of parameters are different in the different theories. Nor is it always evident which of these parameters, or even which of the structural details of the theories, should be regarded as invariant over tasks and which should be subject to adjustment from task to task. Some parts of such a theory are best interpreted as describing fixed characteristics of the organism, others as describing its adoption of strategies to deal with particular task environments. No simple rules will draw boundaries infallibly between invariants and strategies (Gilmartin, Newell, and Simon, 19XX).

Second problem: there exists no defensible body of statistical theory that prescribes how comparisons shall be made (Savage, 19XX; Grant, 19XX; Gregg and Simon, 19XX). Why not

simply prefer the theory that explains more of the variance? If the theories have different numbers of parameters (and especially if the numbers of parameters can't really be counted), comparing the unnormalized variances explained has no intellectual basis.

Tests of statistical significance are even less appropriate here (see references just cited). Moreover, with even a few adjustable parameters, theories can often explain the bulk of the variance in small sets of data -- and the data in the specific situations we shall examine are not voluminous, nor usually reproducible to more than one or two significant figures.

Even if most of the variance has been accounted for, we cannot be sure, without further analysis, what parts of the theory are essential to the good fit, and what parts could be removed or altered without damage to the fit. Sensitivity analysis may show that many of the mechanisms incorporated in the theory have not been tested at all -- that most of the "action" resides in one or a few other mechanisms (Feigenbaum and Simon, 1984).

The testing of theories of the kinds we are examining here calls for far more sophistication than a blind application (or *any* application) of standard numerical statistical tests. It requires detailed qualitative analysis of the relations between data and the specific mechanisms postulated by the theory. Such analysis may even show that, in a case where two apparently different theories explain the same phenomenon, they do so because they possess in common the specific mechanisms that are crucial to these particular phenomena. In the specific respects that are relevant to the comparison, the two theories may not differ at all; they may be essentially identical.

The outcome of a sophisticated analysis will often not be a choice among theories. Rather, it is likely to be the identification of mechanisms, that is, components of theories, that appear to be crucial to the performance. Theories, in turn, can be assembled as various organized combinations of the critical mechanisms. Progress in cognitive science is much more likely to be facilitated by this kind of careful, detailed qualitative analysis than by horse races or formal combat between labeled and trademarked "schools." It is the former kind of analysis we shall undertake in this paper.

Downplaying the role of monolithic theories does not mean that comprehensive computational models of cognition are otiose or dispensable. Nothing will be gained by a return to numerous micromechanisms, each of which explains a different microphenomenon (Newell, 19XX). We need comprehensive theories to permit detailed modeling of phenomena and

comparison over wide ranges of different phenomena. But this does not mean that theories must stand or fall as units, or that they must be treated as black boxes whose contents are not open to examination. On the contrary, progress in theory building calls for the examination of comprehensive theories in utmost detail. This paper essays no more than a very beginning of the undertaking: an examination of two theories in the context of a small range of perceptual tasks.

Specifically, Barsalou and Bower (1984, p. 19) claimed that only parallel-processing models, such as the interactive activation model of context effects in letter perception, could explain pattern recognition:

If several properties of a pattern must be considered during recognition, they are processed serially [by EPAM], so that processing time will increase with the number of tests performed. This prediction clearly runs counter to much of the work in pattern recognition, which generally assumes and has found evidence for parallel processing of stimulus properties (e.g., parallel processing of letters in word recognition: McClelland & Rumelhart, 1981).

In this study we will test whether EPAM, a serial processor, can explain human pattern perception of letters in tachistoscopically presented four-letter words.

Assumptions of the Two Models

The Interactive Activation Model (IAM, McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982) models letter perception in words as an example of the interaction of knowledge and perception. It adds several new features, including interactive processing and the use of inhibition, to the PANDEMONIUM family of parallel processing models.

EPAM (Elementary Perceiver and Memorizer), originally built by Feigenbaum in 1959, is one of the first computer models of human information processing. EPAM has been revised several times, the most recent version is called EPAM III (Simon & Feigenbaum, 1964; Feigenbaum & Simon, 1984). The core of EPAM is a discrimination net and the learning processes that change the net. EPAM III incorporates two learning processes, DISCRIMINATE and FAMILIARIZE. DISCRIMINATE permits a new chunk to be memorized (added to the net) through noting a difference between that chunk and the chunks that are already in memory. FAMILIARIZE acquires information about a chunk that is about to be memorized or is already memorized.

Each of the two models we are considering makes assumptions about the nature of human information processing. IAM assumes:

1. Processing is a multi-level phenomenon. Words are processed at the level of visual features, at the level of letters, and at the level of whole words.

2. Processing is parallel in two ways. All four letters in a four-letter word are processed simultaneously at all three levels and with continual communication between the levels.
3. Interactions between sources of knowledge take place through excitatory and inhibitory activations.

The EPAM model rests on five essential assumptions (Gregg & Simon, 1967).

1. That the human central processing mechanism operates serially.
2. That chunks are the largest stimulus components that are familiar units.
3. That the learning of an item requires a definite amount of processing time per chunk.
4. That human immediate memory has a capacity of a few chunks.
5. That a person can learn any part of a stimulus unit to which he or she attends (i.e., holds in immediate memory for a sufficient time), and that attention may be modified by experimental instructions, attention-directing stimuli, habit, and strategies.

The assumptions of the two models differ markedly in only one major respect: EPAM assumes that the central processing mechanism operates serially, while IAM assumes that processing occurs in parallel fashion. The serial processing assumption enabled EPAM to explain serial-position effects in rote learning experiments and childrens' spelling. It also enabled EPAM to predict that a constant time would be required to memorize a chunk, a prediction confirmed by Bugelski (1962).

EPAM does not assume that every aspect of human processing is serial, only the central aspects involving attention. If word recognition is conceived as beginning at the periphery of the nervous system (the retina) and ending in the central part of the brain (when a pointer to the recognized stimulus is placed in short-term memory), EPAM permits the peripheral processes to occur in parallel fashion, but assumes that the more central parts operate serially. The precise boundary between the parallel and serial processing is not specified.

EPAM shares the IAM assumption that processing occurs at different levels (e.g., word and letter levels) at about the same time, but the two models implement this assumption differently. While IAM stops all processing periodically while the different levels (word, letter, and feature) pass information to each other, EPAM uses the following recursive routine in order to recognize a word:

1. The word recognition process begins at the word level.
2. Whenever a letter must be identified in order to continue the process, the word level suspends activity awaiting the outcome of the letter recognition process for that particular letter.
3. Once the letter has been identified, the word recognition process resumes.

For a more complete description of EPAM III, see Feigenbaum & Simon (1984). For a detailed description of EPAM IV, the revision of EPAM implemented in this study, see Appendix A.

Visual Recognition

For a proper understanding of the EPAM and IAM simulations of word recognition, we must take into account what is known about human visual processing. Neither model has "eyes"; each describes only the central aspects of the perceptual processes. For example, EPAM takes what it calls "objects" as input. An object is a list of letter features, letters, or words. For the simulation, visual stimuli must be represented as objects, and inputting these objects to EPAM corresponds to its "seeing" a word presented in a tachistoscopic visual perception experiment.

The literature on visual perception generally holds that visual images are projected onto an "iconic store" or "visual sensory store," perhaps located in the visual cortex of the occipital lobe of the brain. Generally, the topography of the visual cortex maps the retina, but combines inputs from both eyes. The central portion maps the fovea, representing a visual field of about 0.1 degrees in angle; the remainder maps the periphery of vision, a field subtending an arc of about 3 degrees.

According to Breitmeyer (1984), information arrives at the visual cortex from the retina at various rates of speed. Information about movement in the visual field arrives fastest (through "transient" pathways), and information about contours more slowly (through "sustained" pathways). When a new stimulus activates transient pathways, the information appears to erase the image on the visual cortex in preparation for the more slowly arriving information about the new image that arrives on the sustained pathways. Dow (1974) calculated that the information through transient pathways arrives at the cortex about 1/20th of a second before the information from the sustained pathways. Sometimes the transient information from a new image erases the previous image whose transmission via the sustained pathways is not yet complete, so that the subject never "sees" the earlier image. This phenomenon is called metacontrast.

The icon, which may or may not be located in the visual cortex, has been investigated phenomenologically ever since its discovery over thirty years ago by Sperling (1957). Its salient aspect is its tendency to combine images arriving within 1/10th of a second of each other (Ganz, 1975), and to compare images not separated by a mask when the second follows the first at the

same location on the retina (Phillips, 1974). Just as the visual cortex combines the images on the two eyes, so the icon combines images that arrive within 1/10th of a second (Eriksen & Collins 1967).

Tachistoscopic Perception

Perception is called tachistoscopic when the stimuli are presented for very short periods of time. Most experiments on tachistoscopic perception of words follow the test word with a visual noise mask. A mask is a word-like pattern that is superimposed over the test word. In order to be effective a mask must have contours, which are borders between light and dark. The nervous system appears to use contours in order to identify letter features.

If a word is presented for less than about 20 msec it is generally not seen consciously. We have seen that for a period of about 100 msec after seeing something, everything in the fovea is chunked into a single unit (sometimes called the "unit impulse response"). Generally, in tachistoscopic perception experiments with 4-letter words, the visual noise mask follows the presentation of the word by about 30 to 50 msec. The actual time interval between test and mask is usually adjusted to each subject so as to obtain the same percentage of correct responses.

According to Ganz (1975), there are two prevailing theories to explain why the mask disrupts perception of the word. Both locate the mask's effect in the icon, rather than in the retina, because the masking is effective when the test and mask are presented to different eyes. Masks not composed of visual noise (e.g., a screen of a solid dark color) are only effective when test and mask are seen by the same eye.

One theory of masking, which Ganz calls the "processing interruption theory," assumes that the mask replaces the test in the iconic representation, limiting the time for processing the iconic trace (Sperling, 1963; Averback & Sperling, 1960). However, this theory cannot explain masking where the mask precedes the test word or where the time interval between presentation of the test and the mask is zero or near zero.

According to the alternative theory, which Ganz calls "Integration theory," the contours of the mask are added to the contours of the test stimulus within the iconic representation, so that processing of the icon is disrupted by spurious letter features. This theory has been strongly supported by experiments that show the mask to be more disruptive the more letter-like it is.

The Word-Letter Effect

The word-letter effect was first reported by Cattell (1886), who found that letters presented in words could be identified more accurately than letters presented in meaningless strings. Cattell supposed that words as wholes have properties (like shape) that aid in word recognition. Some later researchers attributed the effect to post-perceptual inferencing.

Reicher (1969) designed a forced-choice experiment that ruled out both of these explanations yet still produced the effect. He ruled out word-shape discrimination by using only capital letters. He ruled out post-perceptual inferencing by presenting choices between alternative letters that, inserted in the appropriate position, all made words.

Reicher found the word-letter effect even in comparing identification of a single letter in a four-letter word with identification of the letter presented in isolation. In his experiment, the single letters could occur in any of eight positions within a two-by-four grid, while the words could be in either of two positions. He suggested several possible explanations, including, this one (Reicher, 1969, p. 280):

A second possibility is suggested by the reports of some Ss that a single letter was harder to find in the field of the tachistoscope than four letters. If the process of perception can be broken down into detection and recognition with the completion of the former necessary before proceeding to the latter, the superior performance on words could be explained in terms of their increased detectability due to the greater area taken up by words than by letters.

Subsequent to Reicher's work, a host of related studies specified the conditions under which the word-letter effect would occur. Two of these studies, conducted by Johnston (1974, 1978), provide a rich description of human performance, and a fertile field for computer simulation. These data were used by Rumelhart and McClelland, and we will use them here to compare EPAM IV's performance with that of IAM.

The Visual Interface

In order to simulate the Johnston data with EPAM, we must specify how the visual stimuli are represented as inputs to the program. We make two assumptions:

1. *Icon as information collector.* The icon collects and encodes parallel information about the visual scene so that it can be recognized by a serial information processing mechanism like EPAM. The encoding is in terms of presence or absence of stimulus features. This latter assumption is also made by the IAM simulation.

2. *Unclassifiable letter features.* Letter features are either seen as present, absent, or

unclassifiable. As a result of the patterned mask, some fraction of the features, distributed randomly, are unclassifiable. The same assumption was made by Rumelhart and McClelland when they used the IAM model to simulate perception under blank mask conditions.

Icon as Information Collector

The function of the icon has not yet been clearly identified. Neisser (1967) supposed that the icon permitted the reconciliation of the visual images in two succeeding eye fixations. However, Phillips (1974), following up upon a study by Phillips and Baddeley (1971), demonstrated that the icon is linked to retinal coordinates and could not reconcile images once the eye had moved. He demonstrated the existence of an alternative visual short-term memory with less capacity than the icon that could reconcile succeeding eye fixations.

The icon collects information from about 100 msec of visual input, a fact that accounts for the sequence of still pictures on a movie screen being perceived as continuous motion.

Some theorists postulate a dynamically changing icon and explain serial position effects as the result of changes in the icon before a given letter is examined. EPAM explains serial position effects as a result of its noticing order: EPAM will terminate its examination of a word when it finds that it cannot recognize the word, and then will only report the letters it has already examined. This produces a serial position curve, for the letters that are examined first have a greater likelihood of being identified correctly.

Unclassifiable Letter Features

We assume that spurious features of the mask within the icon disrupt the features of the target, making those features unclassifiable (rather than incorrect).

For the purposes of these simulations we used the same representation of the letters that was used in Rumelhart and McClelland's IAM simulation. The letters were represented as capital letters in a letter font earlier used by Rumelhart (1970), which classified all capital letters on the basis of presence or absence of fourteen line segments. In order to simulate the tachistoscopic perception task, we gave each feature a .92 probability of being identifiable. At the letter level, EPAM examines about four or five letter features in order to identify a particular letter (26 is nearly 2^5), so that each letter has the probability of about .6 of being identified successfully.

If EPAM were to make the alternative interpretation -- that the masked icon often provided spurious information about a letter, rather than making it unclassifiable, then letters presented in isolation but covered by a patterned mask would often be reported as whole words. To the best

of our knowledge, there has been no report of such a phenomenon

The EPAM IV Model

In this paper, we compare the human data with the outputs of EPAM IV simulations. An earlier version of the program, EPAM III, had been used in previous studies (Feigenbaum & Simon, 1963; Gregg & Simon, 1967). We expected the changes introduced into EPAM IV to enable EPAM to provide a better account of perceptual phenomena, but the innovations can be motivated quite independently of the Johnston experiments. The main innovations were the addition of redundancy in the network of tests and of ability to anticipate characteristics of the stimulus.

Redundancy

In any discrimination net, redundancy permits stimuli to be sorted when information is incomplete, as when the stimuli are viewed from different vantage points or seen under different circumstances. Redundancy is easily added to EPAM III without major changes. With it, EPAM can continue to sort at the word level even when a letter in the word cannot be identified. Two sorts of redundant pathways were added: those to be used if the first letter in the word cannot be classified, and those to be used if the last letter cannot be classified. These pathways make use of the NEC ("not elsewhere classified") branch at test nodes which was already a part of EPAM III. We assume that redundant pathways will develop when the system learns words in situations where the first letter or the last letter is not clearly visible.

Ability to Anticipate

The idea of adding the ability to anticipate characteristics of the stimulus was inspired by CYRUS (Kolodner, 1984), a discrimination net system that built episodic memories. CYRUS developed norms about what to expect within a given episodic category. These norms were built and refined as new episodes were entered in a particular category, and they were added to the information that was maintained for that category.

For example, when CYRUS first developed an episodic category for "diplomatic meetings", in a simulation of Cyrus Vance's memory as Secretary of State, it noticed that in all of the meetings the participants were foreign diplomats and the topic was Arab-Israeli peace, so it anticipated that when Cyrus Vance would meet with foreign diplomats, the topic would always be Arab-Israeli peace. In this case, the anticipations were faulty; Cyrus Vance went on to have

meetings with foreign diplomats about other topics

In EPAM IV, ANTICIPATE adds assumptions to test nodes in the word net. The assumptions apply to the letter position that is tested at the word net node in question. For example, if a node that tests for second letter of a word only differentiates "mare," "more," and "mere," ANTICIPATE will create an assumption list which will include the information that the second letter will be a vowel. This assumption list will facilitate recognition when one of these three words is being read.

In order to implement ANTICIPATE, two additional changes had to be incorporated in EPAM IV. First, a dictionary had to be added, containing information about the properties and features of the letters. Second, two separate recognition strategies had to be provided, one for rapid reading, the other for slow, careful, reading. The rapid reading routine uses assumptions; the slow routine ignores them. The latter is used when the system realizes it has misidentified a word and is now preparing to learn the misidentified word.

The assumptions created by the ANTICIPATE process allow top-down information about words to enter into the bottom-up process of identifying letters. When the letter net is about to sort a letter, the word net will provide assumptions such as whether to anticipate a vowel or whether to anticipate a particular value for a letter feature. These assumptions enable the letter net to sort letters more accurately when a mask has degraded the information in the iconic representation. Hence, they increase the probability of recognizing letters in a word if other letters in the word have been recognized already.

The Johnston Data

Both of Johnston's experiments were reported in Johnston, 1978; the first experiment is described more extensively in his unpublished dissertation (Johnston, 1974). His data measure (1) effectiveness of constraints on letters, (2) a word-letter effect, (3) a word-unrelated-letters effect, (4) a serial position effect, (5) a similarity effect, and (6) a word frequency effect. His data are richer than those of most forced-choice experiments partly because he included a condition where the subjects reported the letters they had seen before picking between forced-choice alternatives. The results (adjusted for guessing) did not discriminate between the condition where subjects reported the letters freely and the forced-choice condition where they did not. Johnston also found that the free reports explain most of the information obtained from the

forced-choice data.

Ineffectiveness of Letter Constraints

Johnston's first experiment produced a "non-effect" that seemed to deny any effect upon letter perception of presenting the letters within a word context. It had been expected that subjects might make sophisticated guesses, supplementing information obtained from the letter itself with contextual information obtained from other letters in the word. Johnston's study appeared to disprove the sophisticated guessing hypothesis, for it showed that people's perceptions of a given letter were not more accurate when there were only two alternatives in that letter position (high constraint condition) than when there were six (low constraint condition). For example, the "S" in "SINK," with many competing possibilities (wink, rink, link, kink, ...) was not harder to recognize than the "S" in "SHIP," with only two alternatives (whip, chip).

(INSERT TABLE 1 ABOUT HERE)

In Table 1, Johnston's data are compared with those produced by the EPAM IV and by two versions of the IAM model. All three simulations fit the human data well, showing very little in the way of letter-constraint effects.¹

We consider first the data from IAM's simulation. According to McClelland and Rumelhart (1981, p. 400), IAM responds more to the number of a test word's "friends" (words that share the critical letter with it) than to the number of its "enemies" (words that differ from it only in the critical letter). Even though low-constraint words have more enemies than high-constraint words (by definition) both sets of words used in the Johnston experiment have about the same number of friends. Moreover, the closest "friend" is the word itself, hence "the node for the correct word dominates the activations at the word level and is predominantly responsible for any feedback to the letter level. (McClelland & Rumelhart, 1981, p. 400)." This feedback is, of course, present in both the high-constraint and the low-constraint conditions, hence produces no difference in the word-letter effect in the two conditions.

EPAM also exhibits almost no letter-constraint effect, fitting the data as well as, or slightly better than, IAM. The success of EPAM is due to the interplay of three factors:

¹We are indebted to James McClelland for making the IAM model available to us. The Interactive Activation Model is the model utilized by McClelland and Rumelhart (1981) when they simulated the Johnston Data. The slight discrepancy between the figures that we report and those reported by McClelland and Rumelhart is probably due to a difference between our random sample of the Johnston words and theirs. The Interactive Activation Serial Position Model incorporates the revisions that Rumelhart and McClelland (1982) made to the IAM model in order that the model be able to simulate serial position effects

1. EPAM's top down component slightly favors the high constraint condition
2. More words in the high constraint condition lay outside of the corpus of words (top 98% of the Kucera and Francis count) that were used by the IAM and the EPAM-IV models, favoring the low constraint condition.
3. EPAM's sensitivity to the informativeness of a letter position in a word favors the low-constraint condition

The third factor is not shared by the IAM model. EPAM tends to examine first those letters that contain the most information, and to ignore those that only contain redundant information. The source of this sensitivity is the DISCRIMINATE process through which EPAM learns new words.

In order to learn a new word, DISCRIMINATE often creates a new test for a letter position in the word net. That test will always be one that differentiates the word being learned from a similar word already in the net. EPAM will not create a test for a letter position that holds redundant information. For example, suppose that the system is creating a test to differentiate among words that begin with *M* and end with *K*, that the word "MARK" is being learned, and that the word "MASK" is the only "M--K" word already known. Then DISCRIMINATE will create a test for the 3rd, rather than the 2nd, letter position, as the second position will not discriminate "MASK" from "MARK." Hence, EPAM will tend to examine important letters before less important letters. If EPAM determines that it can't recognize a word, it will quit trying to recognize its letters; hence letters examined earlier in the process have a better chance of being recognized. Since low- constraint letter positions contain more information than high- constraint positions (they differentiate among more letters), EPAM will generally do a little better with low-constraint positions.

Neither model produces a letter-constraint effect. What the simulations show is that, in both systems, embedding a letter in a word context may prevent the systems from guessing a letter that could not appear in that context, but it does not bias them significantly toward selecting the correct alternative from among those that remain. The mechanisms of EPAM and IAM do not support sophisticated guessing any more than do the strategies employed by Johnston's human subjects. As the next experiment shows, however, the word context in which a letter appears may affect its recognition significantly under conditions different from those of this experiment.

Recognition in Word Context and in Isolation

In his second experiment, Johnston (1978) compared the perception of letters presented in words with the perception of letters presented in isolation, and letters surrounded by unrelated letters. Here, he was replicating the findings of Reicher (1969), who found the word-context effect to occur even where the letters were capitals, and post-perceptual inference was thereby ruled out as an explanation. The words used by Johnston were a sample drawn from the words used in his first experiment; the single letters were shown surrounded by number signs: e.g., "#H##".

(INSERT TABLE 2 ABOUT HERE.)

McClelland and Rumelhart (1981) use IAM to explain the magnitude of the word-context effect under several visual conditions, including some where patterned masks were used. In general, their results fit Johnston's human data closely. According to their model, the reason letters are more accurately perceived in the context of words is because they then "benefit from feedback which can drive them to higher activation levels (page 389)."

The large word-letter effect produced by the IAM model is entirely due to an adjustment by the programmer in one of the IAM model's parameters: When the IAM model is used to simulate perception of a letter within a word or a string of unrelated letters the letter to word excitation parameter is set to 0.07, however when the IAM model is used to simulate perception of a letter surrounded by number signs ("#A##") that parameter is set to 0.00. If this parameter is not changed, the IAM model only results in a 1% word-letter effect (single letters are read with 77% forced choice accuracy) and the Interactive Activation Serial Position Model results in no word-letter effect at all (single letters are read with a 76% forced choice accuracy). For example, if this parameter is not changed, the / in "#l##" is identified more accurately than the / in "FIND" because the many words that have the letter / in their second letter position reinforce perception of the / in "#l##".

EPAM IV also produces a word-context effect, an increase of about 9 in percentage correct. This is smaller than the effect reported by Johnston, an increase of 14, but similar to the increase of 8 reported by Reicher (1967) and the increase of 10 reported by Wheeler (1970).

In Johnston's experiment, the letters are presented surrounded by number signs and followed by a mask. We assume that the letters do not appear in isolation, but as combinations of the features of letters, number signs and masks. We assume that, to recognize single letters

surrounded by number signs. subjects create nets of the form pictured in Figure 1. If they do not recognize a number sign at a particular position in the stimulus, they assume that this position is occupied by a letter and stop processing. A similar net would be created in the situation where a person had to recognize a letter surrounded by blanks, and where the blanks were hard to perceive because of interference from spurious features in a mask. In our EPAM simulations, we collapse these two conditions, using blanks interfered with by masks in place of number signs.

(INSERT FIGURE 1 ABOUT HERE.)

There are two pieces of evidence in the literature to support our assumption that people construct a separate net to sort letters surrounded by number signs, instead of using their letter nets: (1) Mezrich (1973) states that subjects in his forced choice experiment reported using different strategies to perceive letters in words and to perceive individual letters, and (2) Johnston and McClelland (1973) found that leaving out the number signs and using a white (rather than patterned) mask reverses the word-letter effect.

EPAM IV explains the latter finding by assuming that a white mask would not interfere with the perception of a blank. The IAM model does not address directly the finding that letters surrounded by blanks are perceived better than letters in words in the blank mask condition. However, McClelland and Rumelhart do provide an explanation of why, under blank mask conditions, the word-letter effect is smaller for words than for letters surrounded by number signs:

...[When] forced-choice performance on letters in words is compared to performance on letters embedded in a row of number signs (e.g., READ vs. #E##), the number signs serve as a control for lateral facilitation or inhibition. This factor appears to be important under dim-target/blank-mask conditions...

Since there is no patterned mask, the iconic trace presumably persists considerably beyond the offset of the target. It is our assumption that the effect of the blank mask is simply to reduce the contrast of the icon by summing with it. Thus the limit on performance is not so much the amount of time available in which to process the information as it is the quality of the information made available to the system... (p. 389)

Using this explanation their model produces a word-letter effect that is smaller when the letter is surrounded by number signs than the effect in patterned-mask experiments, but larger than the effect in the Johnston and McClelland (1973) experiment.

The EPAM explanation of the word-letter effect in presentation of single letters has three components: (1) EPAM's tendency to sort to most informative letter positions first, (2) EPAM IV's top-down processing to produce assumptions about letters to come, and (3) EPAM's redundant access paths to letters. Of the three, the redundant paths are clearly the strongest mechanism. The word-letter effect would be intensified if more redundancy were added to the EPAM net.

The IAM and EPAM models arrive at similar explanations of the word-letter effect in the patterned mask condition. In both simulations, letters in words have advantages over letters in isolation. In both models the advantages of letters in words are due to the model's use of word information. In both models the subject chooses not to use word information that is available. The explanation of both models, thus, coincides with the explanation of Mezrich (1973) who demonstrated that the word-context effect could be made to disappear if subjects were forced to use the same strategy both with letters in isolation and with letters in words.

In the IAM model, the beneficial effects for letters in words are due to feedback from the word level to the letter level. In EPAM IV, the effects are due largely to the redundancy of the discrimination net for words.

The two models provide entirely different explanations of the effect in the blank mask conditions. EPAM retains the same explanation as in the patterned mask condition, while the authors of the IAM model provide very different explanations in the two conditions.

Words vs Configurations of Unrelated Letters

Johnston's other finding, that letters in words were perceived better than letters in unrelated letter configurations, was also simulated using both EPAM and IAM. Both models obtained the effect using the same conditions that were used for simulating perception of letters in words. With both models, the magnitude of the effect was smaller than in Johnston's human data.

(INSERT TABLE 3 ABOUT HERE.)

Johnston obtained his unrelated letter strings by scrambling the letters in the words by retaining the critical letters in the same positions but scrambling the context letters. Then, since some items remained word-like, he traded some context letters between items. Unfortunately, Johnston did not report the actual unrelated-letter strings used but only reported his procedure for constructing the list, so we had to construct our own list through a similar procedure.

We used a computer to scramble the context letters, then if the string looked like a word, we ordered the computer to scramble them again. In no cases did we switch context letters between items. It is possible that our unrelated letter strings were more word-like than those used by Johnston. See Appendix B for the list of unrelated letter strings that we used.

Serial Position Effects

In his first experiment, Johnston (1974) measured serial position effects for letters in the low-constraint and high-constraint conditions of his first experiment. Since he did not report the actual numbers, we have estimated them from his graphs.

All of the results reported thus far for EPAM IV assume a noticing order 1-4-2-3, that is, the first letter of the word is noticed first, then the last letter, then the second, and then the third. However, people probably do not always recognize letters in the middle parts of words in exactly the same order, hence we also ran EPAM with a 1-4-3-2 noticing order.

EPAM has generally given good quantitative predictions of serial position curves as Feigenbaum and Simon (1984, p. 317) note:

The observed constancy of the serial position curves produced by EPAM can be attributed to the interaction of three features of the architecture: the attention strategy, the anchor point assumption and the approximate constancy of the aggregate set of processes required to learn to recognize and fixate a new chunk in memory. Any system that attended to learning one unit at a time, worked from both ends of the list toward the middle, and required about the same amount of work to acquire each unit would exhibit a serial position curve in quantitative agreement with EPAM's and with those produced by human subjects under usual experimental conditions.

In their second report, Rumelhart and McClelland (1982) manipulated aspects of serial position in patterned mask experiments and then simulated those experiments using IAM. All of their serial position curves were far too flat to match the human data until they revised their parallel processing assumption so that the four letters in a four-letter word were read out at different times. As they noted:

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....

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... 2. It may be that not all letters are read out simultaneously. (p. 76)

Rumelhart and McClelland (1982) then began to alter IAM. First, assuming that letters at the ends of the words were read more clearly than letters toward the center, they changed the four parameters of their system corresponding to the weights of input from the four letter positions. Then, when this assumption could not account for some serial position effects, they changed the read-out cycles for the various serial 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. (p. 77)

When they were still unable to get a close quantitative fit with the serial position data, Rumelhart and McClelland noted that the mechanisms that determine the shape of the serial position curve are "potentially subject to attentional control," and they suggested that "it may be

difficult to gain control and understanding of these mechanisms until the factors that govern control of attention are understood (p. 78).² EPAM does model these mechanisms, with excellent predictions of the human data.

Table 4 and Table 5 show the EPAM data for both noticing orders and for their unweighted average, comparing them with Johnston's human data and with the data from the completely parallel Interactive Activation Model as well as of the Interactive Activation Serial Position Model.² The same data is pictured in Figures 2 and 3.

(INSERT TABLE 4 ABOUT HERE)

(INSERT TABLE 5 ABOUT HERE)

(INSERT FIGURE 2 ABOUT HERE)

(INSERT FIGURE 3 ABOUT HERE)

EPAM achieves a close approximation of the forced choice data, while the Interactive Activation Serial Position Model achieves a close approximation of the free report data. Both EPAM and IAM (in its modified form) explain the serial position effect by means of an anchor-point assumption that treats the first and last letters in a word in a different fashion from the interior letters. Hence, Rumelhart and McClelland's explanation of the serial position effect refutes Barsalou and Bower's (1984) claim that IAM demonstrates that the letters in a word are processed in parallel.

Other Phenomena Simulated by EPAM

In his dissertation, Johnston reported some additional phenomena that were not previously simulated by IAM, but that can be accounted for by EPAM.

Similarity Effect

The "sophisticated guessing theory" tested and refuted in Johnston's first experiment grew out of the observation that mistaken responses in tachistoscopic experiments are often similar to the target words. This similarity effect was present in Johnston's data:

Erroneously reported words do tend to be similar to the stimulus word. On fully 15% of the trials where the correct word was not reported, another word sharing three letters was. While the probability of this type of error due to chance guessing (unguided by any stimulus information) is

²In the Interactive Activation Model, the relative weight parameters for the serial positions were set at 1.0, the mask entered the picture at the beginning of cycle 16, and the letters were read out at the end of cycle 16. In the Interactive Activation Serial Position Model, the positions had relative rate parameters of 1.6, 1.15, .85, and 1.05, the first and fourth letters were read out after 13 cycles, the mask began at 16 cycles, the second letter was read out after 16 cycles, and the fourth letter was read out after 19 cycles.

difficult to estimate, it cannot conceivably be as high as 15%. (Johnston, 1974, p. 136)

When the same statistic is compiled for EPAM IV, in 26% of the trials where the correct word was not reported, another word sharing three letters was. The EPAM IV model thus embodies a similarity effect that is, if anything, too strong. The redundant pathways in EPAM IV, which largely account for the effect, thus also refute a Barsalou and Bower (1984, p. 9) claim that "EPAM nets often do not account correctly for similarity," since such confusions can occur only "when *all* the properties of the stimulus are correctly perceived. (Ibid.)"

We have not calculated this effect with the Interactive Activation Model and we cannot calculate this effect with the Interactive Activation Serial Position Model (Rumelhart and McClelland do not specify when information at the word level is read out in cases where the letters are read out on different cycles), so we are unable to compare the two models with this statistic.

Familiarity Effect

Several researchers have reported that familiar words are more readily identified than unfamiliar words. This familiarity effect is often explained by hypothesizing parallel processing with different "resting" levels of activation for familiar and unfamiliar words in long-term memory. Johnston (1974) noted such a familiarity effect in his data:

Items with a Kucera and Francis-Lorge pooled frequency of 124 or more per million (the upper quartile) had an 8.3% critical letter free- report advantage over times with a frequency of 10 or less per million (the lower quartile). We should expect this difference to be approximately halved in forced choice (since pure guessing would cut in half the number of errors in each set of items), but it actually shrank to only 1.1%. (p. 98)

Johnston does not report which words were included in the high and low frequency categories. Our attempts to reconstruct his words from his definition produced a large set of high frequency words and a minuscule set of low frequency words. Since the EPAM IV and IAM vocabularies were based on the Kucera & Francis word count, we decided to use the Lorge count (Thorndike & Lorge, 1944) to divide Johnston's words into four quartiles. Table 6 shows that when the top quartile was compared to the bottom, EPAM showed a larger high frequency words advantage in both the forced choice and free report conditions than that obtained by Johnston, while the IAM models produced a smaller word frequency effect, much closer to that obtained by Johnston.

(INSERT TABLE 6 ABOUT HERE)

Two aspects of the EPAM IV implementation produced this over-large effect. First, EPAM *always* chose the most frequent word of those accessed by a redundant pathway. For example,

if "CAKE" were the most frequent word ending in "-AKE," then a stimulus sorted to that node would always be interpreted as "CAKE." Second, 15% of Johnston's words did not appear in EPAM's vocabulary because of our decision to use the same vocabulary that was used in the IAM simulations.³

The IAM model comes much closer to simulating the small size of the familiarity effect reported by Johnston partly because it was more successful than EPAM at identifying letters in the 15% of Johnston's words that were not included in its vocabulary. A particular strength of the Interactive Activation Model is its ability to recognize letters in unknown words.

For EPAM IV, a chunk must have been previously learned for it to be recognized. The same is true for the IAM model when reports are made at the word-level (IAM can only report words that are within its memory). However, IAM is able to recognize letters at the letter level in unknown words (and pseudowords) through a process that can be mathematically interpreted as the use of templates (VanLehn, 1984).

We are currently working on an experimental version of EPAM (EPAM V) which uses rule-like templates to translate from one list to another list. The system induces the templates and adds them to NEC nodes as it builds a discrimination net. Although we hope to apply it to other domains as well, EPAM V is initially being used to simulate the reader's ability to recode written words as spoken words. EPAM V first perceives a list by translating it into another list and then translates the other list into a single symbol. It models the process used by people when they first translate written words into spoken words and then translate the spoken words into morphs. Unlike EPAM IV, EPAM V can recognize (by pronouncing) written words and pseudowords that it has never seen previously.

Conclusion

In this paper, we have shown that EPAM IV, a serial recognition and learning system, can account for a considerable number of phenomena that have also been explained by the parallel IAM model of perception. A closer look at the two simulation models showed that their successes have little to do with the question of seriality or parallelism, but are mostly derivable from redundancy, top-down feedback of information from larger to smaller perceptual chunks

³Both EPAM IV and IAM use the vocabulary of four-letter words occurring at least two times per million in the Kucera and Francis (1967) word count. In the Johnston study, 15% of the words do not appear in this vocabulary.

(specifically, from words to letters), and the anchor point assumption.

In spite of the relatively successful performance of EPAM IV, the comparisons of simulations with human data suggest several directions for further revision and improvement of the EPAM model.

First, the top-down feedback in EPAM IV is relatively weak, and contributed in only a minor way to EPAM's explanations of the word- letter effects that were observed in the human data. The EPAM model might be improved by an alternative representation of the interaction between different levels of processing.

Second, the redundancy of EPAM IV, produced too large similarity and familiarity effects. Other methods, such as rule-like templates, should be explored for achieving redundancy.

With these and other directions of exploration open, we can aspire to use the EPAM model in the future to bring about still more complete explanations of human perception of patterns.

References

- Averbach, E. & Sperling, G. Short-term storage of information in vision. *Bell System Technical Journal*, 1961, 40, 309-328.
- Barsalou, L. W., & Bower, G. H. Discrimination nets as psychological models. *Cognitive Science*, 1984, 8, 1-26.
- Breitmeyer, B. G. *Visual masking: An integrative approach*. New York: Oxford University Press, 1984.
- Bugelski, B. R. Presentation time, total time, mediation in paired-associate learning. *Journal of Experimental Psychology*, 1962, 63, 409-412.
- Cattell, J. M. The time taken up by cerebral operations. *Mind*, 1886, 11, 220-242.
- Dow, B. M. Functional classes of cells and their laminar distribution in monkey visual cortex. *J. Neurophysiol.*, 1974, 37, 927-946.
- Eriksen, C. W. & Collins, J. F. Some temporal characteristics of visual pattern perception. *Journal of Experimental Psychology*, 1967, 74, 476-484.
- Feigenbaum, E. A., & Simon, H. A. EPAM-like models of recognition and learning. *Cognitive Science*, 1984, 8, 305-336.
- Ganz, L. Temporal factors in visual perception. In E. C. Carterette & M. P. Friedman (Eds.) *Handbook of Perception Volume V: Seeing*. New York: Academic Press, 1975.
- Gregg, L. W., & Simon, H. A. An information-processing explanation of one-trial and incremental learning. *Journal of Verbal Learning and Verbal Behavior*, 1967, 6, 780-787.
- Johnston, J. C. *The role of contextual constraint in the perception of letters in words*. Unpublished doctoral dissertation. University of Pennsylvania, 1974.
- Johnston, J. C. A test of the sophisticated guessing theory of word perception. *Cognitive Psychology*, 1978, 10, 123-154.
- Johnston, J. C. & McClelland, J. L. Visual factors in word perception. *Perception & Psychophysics*, 1973, 14, 365-370.
- Kolodner, J. L. *Retrieval and organizational strategies in conceptual memory: A computer model*. Hillsdale, NJ: Erlbaum, 1984.
- Kucera, H., & Francis, W. *Computational analysis of present-day American English*. Providence, RI: Brown University Press, 1967.
- McClelland, J. L., & Rumelhart, D. E. An interactive activation model of context effects in

- letter perception. Part 1. An account of basic findings. *Psychological Review*, 1981, 88, 375-407.
- Mezrich, J. J. The word superiority effect in brief visual displays: Elimination by vocalization. *Perception & Psychophysics*, 1973, 13, 45-48.
- Neisser, U. *Cognitive Psychology*. New York: Appleton-Century Crofts, 1967.
- Newell, A., Tonge, F. M., Feigenbaum, E. A., Green, B. F., Jr., & Mealy, G. H. *Information processing language V manual* (2nd ed.). Englewood Cliffs, New Jersey: Prentice-Hall, 1964.
- Phillips, W. A. On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, 1974, 16, 283-290.
- Phillips, W. A. & Baddeley, A. D. Reaction time and short-term visual memory. *Psychonomic Science*, 1971, 22, 73-74.
- Reicher, G. M. Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 1969, 81, 274-280.
- Rumelhart, D. E. A multicomponent theory of the perception of briefly exposed visual displays. *Journal of Mathematical Psychology*, 1970, 7, 191-218.
- Rumelhart, D. E., & McClelland, J. L. An interactive activation model of context effects in letter perception: Part 2. the contextual enhancement effect and some texts and extensions of the model. *Psychological Review*, 1982, 89(1), 60-94.
- Simon, H. A., & Feigenbaum, E. A. An information-processing theory of some effects of similarity, familiarization, and meaningfulness in verbal learning. *Journal of Verbal Learning and Verbal Behavior*, 1964, 3, 385-396.
- Sperling, G. The information available in brief visual presentations. *Psychological Monograph*, 1960, 74, Whole No. 498, 1-29.
- Sperling, G. A model for visual memory tasks. *Human Factors*, 1963, 5, 19-31.
- Thorndike, E. L., & Lorge, I. *The teacher's word book of 30,000 words*. New York: Teachers College Press, 1944.
- VanLehn, K. A critique of the connectionist hypothesis that recognition uses templates and not rules. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, 1984, 74-82.
- Wheeler, D. D. Processes in word recognition. *Cognitive Psychology*, 1970, 1, 59-85.

Appendix A

Description of EPAM IV

EPAM IV was implemented in this study using a version of the IPL-V programming language (Newell et. al., 1964) which was compiled in MSDOS for use on a personal computer. The main aspects of this implementation of EPAM were the representation of the visual input, the interaction of the two nets (the word net and the letter net) in the perception process, the three learning processes (DISCRIMINATE, FAMILIARIZE, and ANTICIPATE) through which the nets were built, and the redundant paths to words in the word net.

Iconic Representation

The "iconic representation" was implemented as a 24 by 5 matrix consisting of locations which can be filled by an "X", a blank, or an "0". An "X" represents a feature that is present, a blank represents a feature that is absent, and a "0" represents a case where the presence or absence of a feature can not be determined. Each letter feature was given an 8% chance of being represented by a "0".

The Word Net

All words in the word net had four letters. The net had a tree structure with a root node at the top where recognition of a word would begin and leaves at the bottom where information about the word would be found. The DISCRIMINATE process which built the net was predisposed to test words for first letter, then last letter, then second letter, then third letter. (An additional implementation was run which was predisposed to the 1-4-3-2 noticing order.) The actual order, however, depended upon the actual words being learned and the order in which they were learned. If a test would be created to distinguish "mild" from "mind", for example, the test would have to examine the third letter position. See Figure 4 for an illustration of a portion of a discrimination net for words.

(INSERT FIGURE 4 ABOUT HERE)

The Letter Net

The root node for this net was a test for type of letter. There were three branches -- vowel, consonant, or "can not determine". The other nodes in the letter net tested for presence or absence of a particular letter feature. The noticing order for this net was determined randomly. One of the letters discriminated had no positive features. This letter was used to represent a non-letter such as the number sign "#", the dash "-" or the blank space. When a letter was

presented in a word-like configuration "##C#", the number sign would be sorted to this node. See Figure 5 for an illustration of a portion of a discrimination net for letters.

(INSERT FIGURE 5 ABOUT HERE)

The Recognition Process

The recognition process sorted a word in the net beginning at the root node and continuing until it either found the leaf that identified the word or until it failed to find a branch at which it could continue. The sorting routine called itself recursively. For example, when the sorting routine was sorting at the root node in the word net and had to determine which branch to utilize, it called itself to sort for the beginning letter position in the letter net. The letter net, then, was sorted beginning at its root node. In order to branch in the letter net, the sorting routine examined specific locations in the "iconic representation". If a "0" were encountered at this location in the "iconic representation", the routine would not be able to continue sorting in the letter net and would report to the word net that it had failed to determine the letter, otherwise the letter sorting process would continue to sort until it reached a letter, at which point it would report the letter to the word sorting process. The word sorting process then would follow the branch based upon the letter found and continue sorting in this fashion until a word was reached. If no letter were found, the word net would follow the NEC branch, if one were available.

Also, the word-sorting process kept track of the letter positions and letters that had been identified so that if sorting would end before a word was found, partial information about the word was still available.

DISCRIMINATE

DISCRIMINATE grew branches so that more leaves could be added to the tree. It worked in two different ways. Sometimes a new test node was added. Other times, a new branch was added to an already existing node. See Figure 6 for an illustration of the cases where a new test node would be added to the net in order to learn the word "MIKE" and a new branch would be added to an existing node in order to learn the word "MIST."

(INSERT FIGURE 6 ABOUT HERE)

FAMILIARIZE

FAMILIARIZE is the process through which information was added to a leaf of the net. This version of EPAM completely familiarizes whenever a word or letter is learned. While complete familiarization would not be appropriate in a simulation of learning, it does speed the

building of the nets in this simulation of perception.

Incidentally, the use of complete familiarization eliminated the natural tendency of EPAM to learn high frequency words before low frequency words. So, in order to counterbalance this effect, words were presented for learning in their order of frequency with high frequency words being presented and learned before low frequency words.

ANTICIPATE

When DISCRIMINATE created a test node or added a new branch to an already existing test node, ANTICIPATE created or refined an "assumption list" at that test node. ANTICIPATE worked differently depending upon whether or not the test node had just been created.

In the case pictured in Figure 5 where a new test node had just been added to the net, ANTICIPATE created a list of those characteristics which the two letters that label branches have in common. In Figure 5, the two letters are "l" and "n". ANTICIPATE consulted a dictionary of the features of letters, and created an "assumption list" which included all of the features, including type (vowel or consonant) that both letters shared in common.

In the case pictured in Figure 5 where a new branch had just been added to an already existing test node, ANTICIPATE refined the "assumption list" at that node by erasing those elements of the list which were not common to the letter labeling the new branch.

For the "assumption list" to help the system with recognition, a new fast-reading recognition routine was added that enabled faster recognition. This routine used essentially the slower recognition process except that when a letter was being sorted in the letter net, the sorting routine would make use of information provided by the "assumption list" at the word net node. Before looking for the letter features of the letter in the "iconic representation," this sorting routine would check to see whether information about that feature was available in the "assumption list." If so, the routine would use that information to enable quicker sorting through the letter net.

The most important information in the "assumption list" was information about the type of letter (consonant or vowel). Where the "assumption list" supplied this information, the root node of the letter list enabled sorting to take a shorter path.

We envision that the rapid routine would be used in most reading situations. However, the slower routine would be necessary when a word would be encountered that was not yet in the EPAM net. Then the system would need to reread without using "assumption lists" so that it could successfully DISCRIMINATE the unknown word and correct mistaken assumptions.

Redundancy

Redundant pathways were included in the word net through the simple process of having the system learn words when the first letter or the last letter could not be classified. For example the word "CAKE" was learned as "CAKE", "?AKE.", and "CAK?". The question mark '?' means that the N.E.C. branch would be followed when sorting this letter. This branch would be used whenever the letter net could not classify the letter in that position. Our assumption was that fluent readers develop the ability to read familiar words even when they don't get a clear look at either the beginning or the ending letters of the word.

In many cases the same redundant pathway might lead to several different words. For example "?AKE" might lead to "CAKE", "RAKE", and "BAKE". In the real world, it is more likely that people would only learn these redundant pathways occasionally, and so it is *more likely* that these redundant pathways would lead to more frequently appearing words. In this simulation, however, in order to simplify the process of compiling our model, we assumed that, in case of conflict, the most familiar word (using the Kucera and Francis count) would *always* be chosen.

Appendix B

Unrelated Letter Strings Utilized

SKNI SPHI WNIK WPIH HLLI HNWE SLLE SENW DTAE DIPR GMAE GIPR DFMA
 DSIK RMPA RSIK BNEO BTAO GNOE GTAO JLLI JNUE TLIL TENU BCKU BMAE TOKU
 TMEA DNAE DCEK PNEA PKCE CTNE CRGA DNET DGRA FTOO FSIT LTOO LSTI FELI
 FEFI WELI WFIE BKEA BTAI WKEA WTAI HVEA HREE WVEA WREE DREA DNAR WREA
 WRAN LTNE LNGO STNE SGNO HOKO HFOO ROKO RFOO SREO SEME TEOR TMEE FDEE
 FRDO WDEE WDRO LAHL DATR LILH DITR SATL SATC TOSL SOTC LALM EANM LILM
 EINM NIDF EINN NODF NOEN KOCN KOMC CUKD CUKM TISM VIEL SOTM EOVL CAKL
 CAEL KICL LIEC SEMS DENM SISM NIMD LELB TELB LOBL TOBL OHWS DHSE WLSO
 DLSE TISM SIKL SUMT SUKL OHTS SHPO SLTO OLPS LILP TIFS LOLP TOSF ALST
 SLRU APST UPRS LATL WAMR LOTL MOWR LAFL TANW LEFL WETN NEDB LELS NIBD
 LISL KACS EABR SOKC EORB SELA ITLH ASNE ITNH EAKF EPKO EFRA EORP EAML
 ANME AEVL AEVN EICD AKCT IEND KINA AGLE PHLE AGME EPMH EANW UHNT AWRE
 TURH OENL EPNT EOSL ETSP NESO AMSH EOTN HATM AECR PTCA RERA PTRR IMLE
 IMLK EMNI MKNI IRSE IWSH EITR HWTI ERLO IPLE EPFO EIPP AELM EHLD AMRE
 DERH KACB FTCA KBSA TASF IENW TTNE IWSE TTSE SBNA SANV ABTS SATV OHNE
 TENV EOSH TESV ELNI AKNR ILCE KACR ALEF LUGF EALP LUGP ADEL AMEL EADN
 EMAN RHAM SETM AHRP TSEP RFOM OOLM RFOK OLOK ARCE OSPE RCAT OPST EBAD
 IARD EABN IARN ESEN ICHN ESEP ICHP ALCP OCRP LCAW ORCW LSIM ITRM LSIP
 RITP EHAL UFOL EHAR OUFR LBAD ANLD BLAE ANLE OBON EHWN OOBT HWET AELD
 OFLD ELAK OLFK INTG RNAG ITNT ANRT AESM ODOM EASR OODR NWID HNAD NIWG
 NHAG RBAK OONK ARBN ONON AEBR EANR EABT ENAT

Figure 1

Word net used to find letter surrounded by blanks or number signs.

1st Letter

#

4th Letter

#

2nd Letter

#

3rd Letter


#

####


Figure 2

Forced-choice serial position curves for Johnston's low-constraint and high-constraint conditions.


Johnston's Human Subject Data



EPAM IV (Combined)



Interactive Activation Model



Interactive Activation Serial Position Model






Figure 3

Free-report serial position curves for Johnston's low constraint and high constraint conditions.


Johnston's Human Subject Data



EPAM IV (Combined)



Interactive Activation Model



Interactive Activation Serial Position Model




Figure 4

Portion of a discrimination net for words.





Figure 5

Portion of a discrimination net for letters.






Table 1

PERCENT CORRECT WITH JOHNSTON'S FIRST EXPERIMENT
HIGH LETTER CONSTRAINT VS. LOW LETTER CONSTRAINT

	Johnston		EPAM IV		Interactive Activation Model		Interactive Activation Serial Position Model	
	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC
Measure:								
Forced- Choice.	80	77	81	78	77	78	75	76
Free- Report of Critical Letters.	54	55	58	56	57	59	52	54
Complete Report of Word.	31	31	35	37				

Table 2

PERCENT CORRECT WITH JOHNSTON'S SECOND EXPERIMENT
WORDS VS. SINGLE LETTERS

	Johnston		EPAM IV		Interactive Activation Model		Interactive Activation Serial Position Model	
	WDS	LTS	WDS	LTS	WDS	LTS	WDS	LTS
Measure:								
Forced- Choice.	85	71	79	70	78	65?	75	66?
Free- Report of Critical Letters.	68	36	57	41	58	35?	53	38?
Complete Report of Word.	43	-	36	-				

Table 3

PERCENT CORRECT WITH JOHNSTON'S SECOND EXPERIMENT
WORDS VS. UNRELATED LETTERS

	Johnston		EPAM IV		Interactive Activation Model		Interactive Activation Serial Position Model	
	WDS	U.L.	WDS	U.L.	WDS	U.L.	WDS	U.L.
Measure:								
Forced- Choice.	85	68	79	71	78	71	75	70
Free- Report of Critical Letters.	68	28	57	42	58	44	53	43
Complete Report of All four Letters.	43	2	36	0				

Table 4

JOHNSTON'S SERIAL POSITION DATA
 PERCENTAGE CORRECT FORCED CHOICE FOR LETTER IN LETTER POSITION
 LOW CONSTRAINT AND HIGH CONSTRAINT WORDS

	JOHNSTON		EPAM IV combined 1-4-2-3 1-4-3-2		EPAM IV 1-4-2-3		EPAM IV 1-4-3-2		Interactive Activation Model		Interactive Activation Serial Position Model	
	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC
POSITION												
FIRST	82	81	82.3	82.3	86	83	78	81	74.9	77.3	78.5	79.1
SECOND	77	76	77.8	74.3	84	83	72	66	78.1	78.4	79.7	80.4
THIRD	76	73	75.7	74.3	71	70	81	78	77.8	78.2	69.9	70.9
FOURTH	81	74	79.9	77.1	81	77	78	77	78.4	79.3	71.9	72.7

Note: The Johnston serial position numbers are estimated from curves drawn in Johnston (1974, p. 85). The actual numbers were not reported.

Table 5

JOHNSTON'S SERIAL POSITION DATA
 PERCENTAGE CORRECT FREE REPORT FOR LETTER IN LETTER POSITION
 LOW CONSTRAINT AND HIGH CONSTRAINT WORDS

POSITION	JOHNSTON		EPAM IV combined 1-4-2-3 1-4-3-2		EPAM IV 1-4-2-3		EPAM IV 1-4-3-2		Interactive Activation Model		Interactive Activation Serial Position Model	
	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC	LoC	HiC
FIRST	56	55	66.0	68.1	74	71	58	65	52.0	56.8	59.4	60.3
SECOND	55	57	55.6	52.1	68	67	43	38	58.3	59.1	61.3	62.6
THIRD	50	49	51.4	50.7	42	44	61	57	57.8	58.5	42.5	44.7
FOURTH	54	53	61.1	63.2	65	63	57	64	59.1	60.7	46.5	48.0

Note: The Johnston serial position numbers are estimated from curves drawn in Johnston (1974, p. 85). The actual numbers were not reported.

Table 6

JOHNSTON'S WORD FAMILIARITY EFFECT
 PERCENTAGE DIFFERENCE IN ACCURACY IN REPORT OF LETTERS
 HIGH FREQUENCY WORDS MINUS LOW FREQUENCY WORDS

	Johnston	EPAM IV	Interactive Activation Model	Interactive Activation Serial Position Model
Measure:				
Free-Report Difference	8.3	14.5	3.9	3.6
Forced-Choice Difference	1.1	9.7	1.9	2.0