

EPAM-like Models of Recognition and Learning*

EDWARD A. FEIGENBAUM

Stanford University

HERBERT A. SIMON

Carnegie-Mellon University

A description is provided of EPAM-III, a theory in the form of a computer program for simulating human verbal learning, along with a summary of the empirical evidence for its validity. Criticisms leveled against the theory in a recent paper by Barsalou and Bower are shown to derive largely from their misconception that EPAM-III employed a binary, rather than n -ary branching discrimination net. It is shown that Barsalou and Bower also failed to understand how the recursive structure of EPAM-III eliminates the need to duplicate test nodes that are used to recognize subobjects, and how the possibility of redundant recognition paths controls the sensitivity of EPAM to noticing order. EPAM is also compared briefly with other theories of human discrimination and discrimination learning, including PANDEMONIUM-like systems and data-flow nets.

Some months ago, the authors of this paper agreed that it was time to review the present state of evidence about the adequacy of EPAM and similar computer programs as theories of the psychological processes involved in recognition and rote learning. The EPAM program is a system capable of learning to recognize and to associate stimuli presented to it, and thus able to simulate the learning and performance behavior of human subjects in verbal learning experiments and related situations. It incorporates a discrimination net as a main component of its recognition mechanism. The recent appearance of a lengthy critique of discrimination nets as psychological models (Barsalou & Bower, 1984) adds to the timeliness of this undertaking.

The present paper is both less than and more than a reply to the paper of Barsalou and Bower. It is less, in that it gives little attention to binary

* The work presented here was supported by the National Institute of Mental Health Grant No. 07722. Correspondence and reprint requests should be sent to Herbert A. Simon, Department of Psychology, Carnegie-Mellon University, Pittsburgh, PA, 15213.

discrimination nets, which occupy about half the space of their critique. EPAM-III (Simon & Feigenbaum, 1964), the version of EPAM that was used in the principal simulations of classical rote-learning experiments, employs general n -ary trees rather than binary trees. This change in structure was accomplished in the years 1961-63, when the authors undertook a general revision of EPAM to remove crudities and limitations of the early versions that had been constructed to test the general feasibility of the approach. Hence, the entire account by Barsalou and Bower of the inadequacies of binary trees comes 20 years too late to be of any particular use or relevance, and will be considered only briefly.

The present paper is more than a reply to Barsalou and Bower for several reasons. First, it is not limited to the discrimination net, which is only one component of the EPAM theory, and by no means adequate, by itself, to predict the psychological phenomena and data that EPAM purports to explain. To understand and evaluate EPAM, it is essential that we discuss the entire structure of the EPAM system, including the mechanisms for control of attention, the learning mechanisms, and the short-term memory assumptions, as well as the structure and behavior of the discrimination net.

Second, the evaluation of EPAM provides a useful vehicle for a more general discussion of what methodologies are best for assessing the adequacy of psychological theories that take the form of computer programs. The Barsalou and Bower paper is very casual and inadequate in this respect. In particular, the authors' frequent reliance on subjective assessments, that certain features of EPAM (or of the imaginary binary-tree EPAM) appear to them to be "psychologically implausible," seems to us a dubious procedure that requires careful scrutiny.

THE ARCHITECTURE OF EPAM-III

The EPAM program undertakes to model a major portion of the system that allows humans to recognize and to learn to recognize stimuli, and thereby to gain access to information related to these stimuli that is stored in long-term memory. We say "a portion," because EPAM represents only a component, although a principal one, of the entire system. According to the EPAM theory, when information is received by the sensory organs (we will have special interest in the eyes and ears), it initially undergoes an encoding process usually referred to as *feature extraction*. The stimulus is then *recognized* on the basis of the features that have been encoded by the feature extraction process, and a symbol is stored in *short-term memory* that "points to" or accesses the relevant information in *semantic long-term memory*.

EPAM models primarily performance and learning in the middle segment of this system: the recognition system and process, the short-term

memory, and the accessing of semantic memory. Therefore, EPAM must be visualized as sandwiched between a sensory-perceptual front end that detects features, and a semantic back end that accumulates and stores, in associative fashion, information about the things that have been discriminated. In the recognition process, EPAM's input is the set of features presented to it by the sensory-perceptual mechanism, and its output is the symbol that accesses an *image* stored in semantic long-term memory. The feature extraction process is not modeled, nor is the whole structure of schemas and associations in long-term memory. EPAM's model of short-term memory is simplistic—a small set of slots, so to speak. Most of the detail of the model is reserved for the discrimination net, the images, and the control structure. The learning component of the theory is also restricted to the modification of the discrimination net and the images.

In short, we think of the semantic long-term memory as the “encyclopedia” in which most of the knowledge acquired by the learning system is stored; the discrimination net as the “index” to this “encyclopedia,” a recognition memory; and the short-term memory as a small working memory needed by learning processes.

Discrimination Net and Recognition Process

EPAM uses a homogeneous representation for the knowledge that it learns and accesses: the object. The same representation is used for external objects (“stimuli”) and internal objects (“images”). An object is either a list of parts (called “subobjects”) or a set of features or properties. The object itself can have properties of its own that are different from the properties of its subobjects. The list of subobjects is an ordered list. This representation is recursively defined, i.e., a subobject has the same representation.

For example, an S-R would be an object consisting of two subobjects, the S followed by the R; and it might have a property (e.g., its color is red). The S might be a syllable of three letters, hence three subobjects. The same might be true of the R. The letters themselves are subobjects, but in EPAM they are represented as “primitive,” i.e., having no subparts but only descriptive features.

During the learning activity, the image is built. It is possible that it contains some but not all of the information present in the stimulus object. Nevertheless, it is represented in the same form as the “external” stimulus object.

The Net. The discrimination net (the “index”) is a branching tree structure. The various (and usually numerous) leaves of the tree are the images (the memorized information). All other nodes of the tree are test-nodes, at which are stored tests to be performed by the recognition process. In EPAM III,

each test-node is an n -ary switch (not the binary switches of EPAM II that proved to be inadequate).

The tests stored at the test-nodes are, in keeping with the basic representation, either subobject tests (s.o. tests) or property/feature tests (p tests). The branches that lead to the next level in the tree are accordingly labeled with either the (internal) name of an image or the value of a property.

Of the n branches at a test-node, one is special.² Its label is "not elsewhere classified at this node" (n.e.c.). During learning, it routinely happens that an object being recognized is "new," is not yet discriminated at the test-node (i.e., doesn't have a labeled branch) or does not even have the s.o. or the p being tested there. For such objects, n.e.c. is the branch-path.

EPAM III makes salient the distinction between "types" and "tokens" of memorized objects (familiar objects). The type is the image of the familiar object and is used during recognition to answer the question, "Do I recognize this object that has been passed through the discrimination net?" But the familiar object may occur in many different symbol-contexts, e.g., the letter "A" may occur in many words or the word "CAT" may occur in many different lists. A token of the familiar object is needed to indicate the presence of the object in its various contexts. The token need not be a complete copy of the type (although that is one possible theory); it can be simply a pointer. Since the type is an image; and the image is a memory structure that has a unique internal symbolic name; that name can serve as the token.

The Recognition process. In the Recognition process, the discrimination net and the image are the two memory structures used to decide whether an object presented is already familiar ("do I recognize this?"). The test-and-branch processing in EPAM III is not entirely straightforward as it was in EPAM II, because the recognition process makes use of the recursive definition of the representation of an object. The recognition proceeds as follows:

The object O is presented. The first test-node is the root node of the discrimination net. If the test is a p test, $v = p(O)$ is determined; v is found in the node's set of labeled branches, thereby accessing the next test-node of the net. (Recall that there can always be an n.e.c. branch). If the test is an s.o. test, the appropriate s.o. of O is accessed, but branching to the next test-node can not immediately be done, because the s.o. must first *itself* be recognized before a search for its appropriate labeled branch takes place. So the process of recognizing O is suspended temporarily while the recognition of s.o. is accomplished. But the recognition process is the same, i.e., recursively homogeneous. S.o. begins its traversal of the net. Further subtruc-

² This is a slight simplification for purposes of exposition. There are, in fact, *two* special branches at each test node. The first is used for stimuli that do not possess the property or subobject being tested; the second, for stimuli that have a new value, not previously encountered, of the property being tested.

ture of s.o. may be encountered, in which case the process recurses yet another level. Eventually the recognition of primitives terminates all levels of recursion (recall that primitives are objects having only p tests). If the image of s.o. is found, its token (internal name) will be the label of the appropriate branch to the next test-node. If not, the n.e.c. branch is taken.

To summarize, the recognition of an O that is not primitive will consist of several subpart recognitions and subsubpart recognitions. There will, in general, be several “open decisions” in suspense awaiting the completion of further subpart recognitions.

For example, to recognize the S-R pair CAT-DOG, the recognition process may first have to recognize the s.o., CAT, as a familiar entity. To do that, it may have to recognize the s.s.o., C, as a familiar letter. The recursion need go no further since C is primitive.

Recognition consists not merely in finding some image (i.e., arriving at some leaf of the tree) but in finding one that matches (to the extent of the information available) the object being passed through the discrimination net. The match processing is done when the recognition process encounters a leaf of the tree (i.e., image stored here, not test). The match will “fail” (i.e., no successful recognition) on CAT versus CAR, for example; but will “succeed” on CAT versus CA— since the CA— may very well be a partially memorized image of CAT. Match failures occur because differences are found between object and image, and these differences can be exploited by learning processes to learn further discriminations, i.e., grow the discrimination net (“expand the index”). We turn our attention next to learning processes.

Learning Processes

EPAM III has two learning processes: FAMILIARIZE and DISCRIMINATE. For convenience in the description, we will abbreviate these as FAM and DISC. FAM is the memorization process by which the images of objects are gradually built up, adding properties and subobject tokens. DISC is the discrimination learning process by which the discrimination net structure is augmented (test-nodes and branches being added) to allow correct recognition as the number of learned items grows. More images, of course, imply more branches are needed, and often more test-nodes.

Familiarization. FAM begins by sorting an object through the net. If no image is found, an initial image is created containing only the information that was examined along the path to the net leaf. However, if an image exists, FAM checks object and image for differences. If no differences are found, at least to the extent of the (possibly incomplete) information in the

image, some additional properties or subobjects are added to the image. Obviously, after a number of such learning events, familiarization will be complete and there will not be any more incremental memorization relevant to that image.

Discrimination. If differences are found, that means a confusion (failure to discriminate adequately) exists, and the discrimination net must be augmented. The strategy of DISC is to modify the discrimination net as little as possible to accomplish the augmentation. The existing net tests are used to the extent possible to accomplish the discrimination. Starting with the root node of the net, the tests are applied to both object and discrepant image. If a test is encountered for which the outcomes differ, which can happen if the image is attached to a n.e.c. branch, the object and the image are each given a new branch appropriately labeled with its test value. Thus, discrimination learning is forced to the highest possible levels of the existing tree before any new subtrees will be added. For an understanding of subsequent sections of this paper, it is important to keep in mind that in the n -ary nets of EPAM III, net growth is biased toward breadth, rather than depth.

For example, suppose the net thus far consists of a test for initial letter and a test for final letter, that at the first test there is a branch for the letter C, but at the second, no branch for the letters T and W. Then CAT and CAW will both be sorted by the first test to its branch C, but by the second test to its n.e.c. branch. Suppose the image stored at that branch is C-T. Then a difference between the final T and W, respectively, will be noticed, and new branches, T and W, added at the second test node. An image like C-T will be stored at the leaf of the former branch and an image like C-W at the leaf of the latter branch.

If no existing test-node can be used to make the needed discrimination, another must be added, i.e., net growth in depth must occur. The list of differences is used as the basis for the growth. In EPAM III, a growth of one test-node per occurrence was used in most experiments, though the number is a parameter of the mode. P-tests are given priority for net growth. If no property differences have been found, then s.o. tests are used.

In the previous example, if COT is now confused with CAT, a new test, for second letter, will now replace the C-T leaf with branches for O and A, and with images of COT and CAT at the leaves of these branches.

Noticing order. The order of use of these discriminating properties and subobjects is important and was extensively studied with the NOTICING ORDER mechanism of the EPAM II model (Feigenbaum, 1959; Feigenbaum, 1961). The noticing order is important because it encodes much of the learner's strategy in the familiarization of new material, and the learner's adaptation to the locus of differences that could be used for discrimination

learning. In EPAM II, the order of noticing differences was made highly adaptive to the stream of new objects presented for learning. This was true of both properties and subobjects.

Suppose a stream of rather heterogeneous three-letter syllables was being presented for learning, and the learner's initial strategy was to discriminate on the basis of tests on the first letter of each. Suppose that the experimenter changed the nature of the stream so that all first letters were the same. EPAM II would rapidly shift its attention to looking directly at third-letters for discriminating differences. The noticing order itself was changed, so that subsequently no processing effort would be expended in the fruitless search for first-letter differences (until such time as attention to third-letters no longer paid off, in which event attention would shift to second-letters or back to first-letters).

The noticing order mechanism of EPAM II was available to be used in EPAM III as the basis for ordering p-tests and s.o.-tests inserted when growing the net in depth. But because this part of the model was not essential to the experiments that we wished to perform with EPAM III, the adaptive noticing order was not used in EPAM III.

FAM and DISC work together; they are not really separable learning processes. FAM must know that the image being augmented is uniquely discriminated. To insure this, DISC must use differences that can only be found if FAM has memorized enough detail in the image. To insure that there will, most of the time, be enough detail for discrimination, FAM is called upon to add detail to the image (do more familiarization) even when the model is making correct responses in the learning experiment, i.e., some "unmotivated" familiarization is allowed to occur.

Semantic memory. The EPAM III model does not address the question of the formation of associations between images and relational clusters of images that cognitive scientists denote with the term "semantic memory" (the "encyclopedia" mentioned earlier). In the post-EPAM years, the issue of modeling semantic memory was addressed by Quillian (1966, 1968) and subsequently by many others (Anderson & Bower, 1973). The intention of the EPAM modeling work was to address the issues of the memorization of objects and the discrimination among them (the "index" issue).

EMPIRICAL EVIDENCE FOR THE EPAM THEORY

If one's only acquaintance with EPAM came from the Barsalou and Bower paper, one might well wonder why a model with so many defects, and so lacking in positive virtues could ever have been advanced as a serious scientific theory. There are at least two answers to this question. First, and as we

shall argue in later sections of this paper, it is debatable, to say the least, as to how defective the claimed "defects" are. Second, apart from one passing reference, Barsalou and Bower do not mention, much less discuss, the rather long list of significant psychological phenomena that are successfully predicted and explained by EPAM, and that have not been equally well predicted and explained by any other theory known to us.

Philosophers of science (we have in mind particularly Popper (1959, 1965) and Lakatos (1970) remind us that in any field of science it is seldom that the reigning theory is not faced with one or more anomalies—facts that do not fit. If the theory has been successful in accounting for a great many phenomena, the presence of anomalies is not usually regarded as a sufficient reason for rejecting it. Instead, anomalies focus new research on ways of improving the theory to remove them, or alternatively, ways of accounting for them, that deprive them of their anomalous status. Rejection waits until a new theory has been proposed that handles the anomalies (or some of them) without losing the predictive and explanatory power of the original theory.

Evaluation of a scientific theory, then, requires at least as much attention to the positive evidence for the theory as it does to anomalies or supposed anomalies. We demand of theories that they tell the truth, but not that they instantly tell the whole truth, or nothing but the truth. If we demanded more, we would have no scientific theories. Let us now consider what positive evidence might lead us to believe that EPAM-III is a good theory for explaining the known phenomena that surround the processes of recognizing familiar stimuli and of learning to recognize and discriminate new stimuli.

Positive Accomplishments of EPAM

The two earliest versions of our theory, EPAM-I and EPAM-II, provided some sufficiency tests of the architecture. They showed that the set of mechanisms incorporated in EPAM was sufficient for discrimination and recognition, and for the learning of new discriminations. These programs did, in fact, learn lists of nonsense syllables. But they were more successful than that, for they also made correct quantitative predictions of a number of phenomena that at just this time (around 1956 to 1962) were receiving attention in the literature on rote serial learning. The principal publications reporting these predictions were Feigenbaum (1961), Feigenbaum and Simon (1961), and Feigenbaum and Simon (1962).

1. They accounted for the constancy of shape of the serial position curve, which had been discussed by McCrary and Hunter (1953). Not only did they predict correctly the "standard" shape of the

curve, but they also explained why the curve deviates from this shape under particular conditions: specifically, when certain items on the list are perceptually salient (von Restorff effect), and when subjects' attentional strategies are modified by instructions. Subsequently, the same explanation was extended in the SPEL program to make very specific and accurate predictions of children's spelling errors (Simon & Simon (1973).

2. They predicted that a constant time would be required to fixate a single item, independently of the drum rotation speed. This constancy could have been (but wasn't) inferred from experiments already in the literature, but was first addressed explicitly (and confirmed) in experiments by Bugelski in 1962.
3. They also postdicted the finding of Rock (1957), unknown to the authors when EPAM I and II were constructed, that under certain circumstances syllables would be learned in a single trial.
4. They exhibited, by forgetting items previously learned, both oscillation (Hull, 1935) and retroactive inhibition, phenomena already familiar from the experimental literature.

We will mention later other correct predictions of EPAM, but it was the four mentioned above, and their ramifications, that provided the main confirmation that EPAM was on the right track.

The two major innovations introduced into EPAM-III were both modifications in the architecture of the discrimination net. The first was to make the structure of the net recursive; the second was to replace the 2-way branches of EPAM-II with n -way branches. The first of these changes is described in some detail in the first published report of the results of running EPAM-III (Simon & Feigenbaum, 1964); the second is mentioned only glancingly in that paper, and was first made fully explicit in Newell & Simon (1972, pp. 34–36). In the 1964 report, the authors merely say:

At each node of the net, some characteristic of the stimulus is noticed, and the branch corresponding to that characteristic is followed to the next node.

Similarly, in Feigenbaum (1965), a paper cited by Barsalou and Bower, we find (p. 40):

This memory has a *tree* structure, called the *discrimination* net. At each nonterminal nodal level is stored a testing process, a *discriminator*, which tests some feature of a stimulus object and sorts the object along the appropriate branch to the next nodal level.

Although these are inexplicit ways of referring to an n -ary branched tree, they are certainly not the ways one would refer to a binary tree. The phrase, "the branch corresponding to that characteristic" is not a way of describing a negative test. However, even if the reference to n -ary branches is allusive,

this allusiveness cannot be blamed for Barsalou and Bower's supposing that the binary assumption was present in EPAM-III, and even intrinsic to EPAM-like theories. For in the book, *Human Associative Memory*, that Bower wrote with John Anderson in 1973, the multiple branches of the nodes of EPAM-III were explicitly noted, and even illustrated with a diagram (see Figure 1)! On page 70 of that book, we find:

In practice, the early versions of EPAM and all versions of SAL tested for the presence or absence of a particular feature or attribute-value pair on the description list of a stimulus. Thus, the output of the test was either *yes* or *no*. *A later version of EPAM used tests for the value of a specified attribute, so multiple output branches were possible* (Our italics)




Figure 1. A Schematic Illustration of a Discrimination Net Used in EPAM and SAL (after Anderson & Bower, 1973, p. 70).

Immediately below this paragraph is a figure of a multiple-branch discrimination net captioned: "A schematic illustration of a discrimination net used in EPAM and SAL."

The concerns that led us to introduce recursive structure and n -way nets into EPAM-III included some of the same concerns that Barsalou and Bower rediscovered 20 years later and used as foundation for their criticism of binary nets. One was the elimination of heavy reliance on negative tests; the other was the need to remove the duplication of clusters of subtests when the same substructures reappeared in different contexts as components of larger structures.

With these changes in place, EPAM-III appeared to be ready to confront the main phenomena that had been described in the literature of rote verbal learning, and which would, therefore, have to be explained. In addi-

tion to handling the phenomena already discussed above, EPAM-III and a simplified version called MAPP successfully explained the data in at least four experimental environments that were then or later prominent in the literature (Simon & Feigenbaum, 1964; Gregg & Simon, 1967; Simon & Gilmartin, 1973):

5. EPAM predicted, with good quantitatives as well as qualitative agreement, the effects of familiarity and meaningfulness on the time required to learn serial lists or sets of paired associates. In particular, it explained the ubiquitous upper bound of three to one on the ratio of time required for trigrams of lowest meaningfulness as compared with trigrams of highest meaningfulness or meaningful words. Specifically, learning time was shown to be proportional to the number of new chunks that had to be associated, where "chunk" has exactly the same meaning as in G. A. Miller's (1956) model of short-term memory. Thus, converging definitions were provided that permitted the reality of chunks to be verified (Simon, 1974). The fixation time per chunk proved to be about 8 s.
6. EPAM predicted, with good qualitative agreement, the effects of intralist and interlist stimulus and response familiarity.
7. Meanwhile, although the earlier versions of EPAM had already provided an explanation of the conditions under which learning would take place in a single trial or would be incremental, this explanation had been little noticed, and these phenomena were still regarded as mysterious in 1965. In that year, Atkinson, Bower, and Crothers (1965, p. 118) wrote:

At present writing a challenging task confronting the theorist is (1) to identify those experimental conditions where the one-element (i.e., all-or-none) model works, and those where it does not, and (2) to construct a more general model which essentially reduces to the one-element model in the former cases but accounts for the discrepancies in the latter cases.

Gregg and Simon (1967) responded to that challenge by showing that EPAM-III could account for the phenomena, with good qualitative and quantitative agreement with the data.

8. Using a program called PERCEIVER, Simon and Barenfeld (1969) successfully simulated the sequence of eye movements of an expert chessplayer looking at a board position. In their report of this simulation, the authors observe that for a more comprehensive explanation of the perceptual and memory capabilities of chess experts and masters, the PERCEIVER mechanisms would have to be complemented by an EPAM-like system capable of recognizing many

thousands of patterns of pieces on the board. Subsequently, Simon and Gilmartin (1973) constructed the MAPP program, which adds such an EPAM-like component to PERCEIVER, and showed that it could simulate the experimental evidence on expert recognition of chess patterns that had been accumulated by de Groot (1965) and by Chase and Simon (1973). The discrimination net used by Simon and Gilmartin (1973) was less sophisticated than that of EPAM-III, and employed binary tests; an n -ary version was later developed but not reported in the published literature. The distinction between n -ary and binary branches was not relevant to the particular phenomena under study.

These eight items do not exhaust the set of psychological phenomena for which EPAM makes correct predictions. For example, Brown and McNeill (1966) showed that letter-order effects in the so-called "tip of the tongue" phenomenon were just what is predicted by the EPAM model. We will refer to other experiments providing positive support for EPAM as we proceed.

This list of findings represents an impressive series of successes for the EPAM theory, especially in the form in which it was formalized in EPAM-III. We cannot point to any other theory in cognitive psychology, whether expressed verbally or mathematically, that has accounted for such a wide range of empirical data, or that has explained so many phenomena to which the literature had called attention as needing explanation. This record of success does not put EPAM beyond criticism, but it surely establishes its claims to serious consideration. As we learn more about the phenomena of learning and recognition, the theory will obviously be modified and extended. Features of it may well prove to be simply wrong. But it would be very surprising, in the light of this record of success, if an important measure of truth were not found at its core.

Attribution of EPAM's Successes to Component Processes

In the remainder of this paper, we will be examining the criticisms that have been leveled against the EPAM theory, assessing the validity of the arguments and the evidence on which they are based. In each case where EPAM's behavior is questionable in the light of the evidence, we will wish to trace the difficulties back to particular features of the system so as to see more clearly how it should be modified. All-or-none verdicts will provide us with little help in developing further the theory of recognition and learning.

As a first step in carrying out such a critique, it will be useful to make a similar assessment of EPAM's successes. Not every instance of correspondence between data and EPAM's behavior provides supporting evidence for

all of the features of EPAM. Particular experimental results may be due to particular individual features of the simulation system. We will comment, in turn, on each of the sets of findings that were listed before.

1. *The serial position curve.* 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.

The data on the serial position curve do not test details of the structure of the discrimination net. EPAM-II, with its binary net, produces about the same curve as EPAM-III, with its n -ary net. The data do test the other three aspects of the theory that we have mentioned.

2. *Constant fixation time.* All versions of EPAM predict, correctly, that total fixation time will be independent of memory drum speed (within a wide range of speeds). This follows from the fact that the amount of fixation that takes place is a function of the amount of time available for carrying out the net-growing and image-storing processes that are required for fixation. Total fixation time is independent of the number of trials or the time of exposure per trial. Again, this result is not sensitive to details of the architecture of the discrimination nets.

3. *One-trial versus continuous learning.* The appearance of one-trial learning depends upon the simplicity or complexity of the learning required for each item, and upon the attention strategies adopted by subjects. The structure of the discrimination net determines what constitutes a "simple" or a "complex" item, and thus is implicated in this phenomenon. It would not be easy to produce the correct behavior in a system that did not build up its stimulus units into recursive structures of hierarchically organized chunks.

4. *Oscillation and retroactive inhibition.* Forgetting of items previously learned was produced by the gradual modification of the discrimination net during learning so that items previously sorted correctly would now be sorted to the wrong image. To produce this effect, in the course of the learning process a recognition system would have to change continually the connections between the feature tests and the memory images of the items already recognized. This phenomenon might occur in recognition structures other

than discrimination nets, but it cannot be taken for granted. It would have to be demonstrated for any structure that is proposed to replace discrimination nets.

5. *Effects of familiarity and meaningfulness.* The prediction that meaningful words will be fixated in about one-third the time required for nonsense syllables of low meaningfulness depends again upon the recursive structure of the chunks into which stimuli are organized by the system, and this depends, in turn on the architecture of the discrimination net and its processes.

6. *Effects of similarity.* The effects of similarity in an EPAM-like structure depend very sensitively upon the distance or proximity of items in the discrimination net. Hence, experiments that manipulate similarity provide a powerful means for evaluating the fine structure of the recognition system.

7. *One-trial learning.* This topic has already been discussed under Item 3.

8. *Expert pattern perception.* These experiments provide further tests of the system's capabilities for acquiring patterns and for recognizing instances of the patterns it has acquired. They do not provide sensitive tests of the detailed architecture of the discrimination system.

This brief review confirms the necessity of viewing a system like EPAM, not as an unanalyzed complex black box, but as an organized system of components and assumptions, each of which requires its own evidentiary support. We see that the experiments with EPAM have provided some support for each of its major features.

TESTING A THEORY'S VALIDITY

Of course there is no question of "validating," once and for all, a theory that makes universally quantified assertions, as scientific theories do. Conclusive validation is impossible for two reasons: first, because there can be no guaranty that a new observation will not refute a previously supported theory. Second, there is no guaranty that the theory provides a *unique* explanation of the phenomena.

Falsification of Theories

Even if theories cannot be conclusively validated, perhaps they can be conclusively falsified. In the crystal-clear, but rarified, atmosphere of logic, a single contradicting fact will falsify a theory that is universally quantified. Even one A that is not B falsifies the proposition that all A are B.

In the real world of science, falsification is no simple matter. For the interpretation of any experiment or any observation hinges not only on the theory or theories that are being tested. It depends also on a whole host of auxiliary assumptions that are implicit (or explicit) in the methods and instruments of observation, and in the processes for deducing the implications of the observations for the theory under test. These auxiliary assumptions have been called the “protective belt” of a theory (Lakatos, 1970), for they provide alternative interpretations of an apparently falsifying observation, interpretations that may save the theory.

What an apparently falsifying observation normally initiates is not an immediate rejection of the theory but a long and complex process of trying to diagnose and repair the difficulty, preferably by localizing the trouble in the theory’s auxiliary assumptions or protective belt. Meanwhile, the falsifying observation is treated as an anomaly—a disfiguring wart on the face of the theory that may or may not turn out to be malignant.

If an alternative theory is available, we may be more willing, in the face of anomalies, to give up the one that has reigned thus far. And that is why we pay particular attention to the critical experiment—whose interpretation is no more immune than the interpretation of any other experiment from the possibilities of manipulating auxiliary assumptions. In cases where the current theory has been highly successful in some domain (e.g., Newtonian mechanics), it is sometimes saved as the limiting case of the more successful theory that replaces it (i.e., relativity theory).

Thus, in the natural history of science, theories are not only born and flourish, but they also evolve as new evidence forces their modification. And they fight fiercely for their existence against competitors that try to crowd them out. Death, for theories as for organisms, is only the final failure in this struggle for survival.

“Implausability” and “Inefficiency” Arguments

Barsalou and Bower, in their article, employ three kinds of evidence as a basis for their skepticism about discrimination nets in general and EPAM in particular. They cite a number of experimental studies which they allege falsify the theory, or some aspect of it. They describe certain components of the theory as “psychologically implausible.” And they reject certain features of the theory as unacceptably inefficient. We have already considered the nature of falsifying experimental evidence. We now need to make some remarks about the arguments of “psychological implausibility” and “unacceptable inefficiency.”

To say that something is implausible implies that we already have some presuppositions about how it must be. To say “it is implausible” is equivalent to saying “I have an intuition that . . .” Now the intuitions of experts

are not to be sneezed or sneered at; they commonly have at their origins a good evidential base. They are "intuitions" because the expert can confidently reach the conclusion without reviewing or recalling the evidence (Chase & Simon, 1973). However, when one is trying to make one's intuitions plausible or convincing to another, that is just what must be done—the evidential base must be recovered.

For example, we might think it "psychologically implausible" that the recognition process is a serial rather than a parallel process. We may no longer be aware of the evidence that led us to this conclusion, for the issue may have been long laid to rest in our minds. When memories of the evidence are stirred again, we may find that the belief in implausibility stems from the fact that recognition takes place very rapidly, or from our knowledge of physiological evidence that the retina and association areas show extensive parallel electrical activity.

Hence, the feeling of "psychological implausibility" can be the beginning of the analysis, but it cannot be the end, unless we are in the presence of persons all of whom share our intuitions. We cannot settle an issue like "serial versus parallel" by arguments of plausibility unless our listeners are already convinced. We must go back to the empirical evidence that produced this plausibility in the first place.

The argument from "unacceptable inefficiency" is another matter. It is essentially a Darwinian argument: that in the evolutionary process only efficient systems survive. Sometimes the argument is even strengthened to read: in the evolutionary process only optimally efficient systems survive.

The stronger argument—the argument for optimality—is patently invalid. At most, evolutionary processes strive toward local maxima; they provide no guarantees that global maxima will be attained. Nor does a system have to maximize anything to survive; it simply has to be more efficient than its competitors. The empirical evidence of evolution bears out this simple truth. When islands like Hawaii and Australia are invaded by creatures from which the oceans have long protected them, the new immigrants frequently thrive and drive out native species. If optimization had taken place previously, that could not happen. It can only occur if the immigrant is more efficient than the less-than-optimal natives.

Consequently, we must interpret the argument of "unacceptable inefficiency" in its more moderate form. Proposals of very inefficient mechanisms must be rejected, where "very inefficient" means less efficient than other mechanisms *which surely can be counted upon to have evolved by now*. But the final clause, which cannot be omitted, deprives the inefficiency argument of almost all its force. Natural systems are full of all sorts of arcane and baroque subsystems for which a clever engineer might well devise simpler and more efficient replacements. There is no reason to believe that the brain, the newer parts of which have only been subject to 10,000 years

of selective evolution, is designed in any spectacularly efficient way. As Dr. Johnson said of the dancing dog: "The marvel is not that it dances well, but that it dances at all."

Survey of Evidence Against EPAM

Let us now examine the ways in which Barsalou and Bower support the arguments they make against the EPAM theory. We found in their paper 35 instances where they gave reasons for their objections to EPAM as a psychological theory. Of these 35 instances, only 8 (23%) referred to empirical evidence drawn from experimental studies. Even in these 8 cases, the findings of the experiments were used only qualitatively—no quantitative comparisons were made in the entire article between the behavior of EPAM and empirical data. The following is a typical example of a reference by Barsalou and Bower to an empirical study: "Discriminativeness has been shown to be central to human performance on numerous other occasions (cf. Fisher, 1981; Krumhansl, 1982; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Winograd, 1981)" (p. 7).

What is the nature of the remaining "evidence" that Barsalou and Bower adduce? In 13 cases (37% of the total), they condemn features of EPAM as "psychologically implausible." Two examples should suffice to illustrate this kind of argumentation: "Second, it seems counter-intuitive that people typically recognize things using negative information" (p. 4). or "But such proposals, however efficient they may eventually be, seem psychologically implausible" (p. 17).

In 6 cases (17%), Barsalou and Bower simply appeal to authority, as though these issues were to be settled by a majority vote among psychologists. One example will suffice: "Yet it is now widely believed that human pattern recognition is a parallel process" (p. 11). Three sets of psychologists who are well known to be believers in parallel processes are then cited as a kind of Gallop poll to substantiate this claim of majority support.

Finally, in 8 cases (23%), a feature of EPAM is discredited on the grounds that the mechanism is "inefficient." We read, for example: "Obviously this is not an efficient use of storage space" (p. 11), or "When a new domain does not share properties with old domains, this contingency is unnecessarily and inefficiently represented in EPAM nets as a long path of negative properties" (p. 14). (This claim happens to be incorrect, but for the moment, we are not concerned with its accuracy but with its evidential basis.)

Thus, when we look more closely at the claims of Barsalou and Bower, we see that they are just that—claims—and that the evidential basis for them is hardly commensurate with the space required to state them. Moreover, at least 12 of the 35 arguments made against EPAM disappear as soon as we

dismiss the false notion that the EPAM nets are binary. In the next sections we will see that there is considerably less than meets the eye in the remaining 23 objections.

NEGATIVE PROPERTIES

We shall take up the objections to EPAM in the order in which they are considered by Barsalou and Bower, using their categories for them. Their first problem with EPAM nets, "their heavy reliance on negative properties" (p. 4), can be treated briefly since the objection rests on the misconception that EPAM nets are binary. As Barsalou and Bower themselves acknowledge (p. 17), once n -way branches are admitted, the difficulties with negative properties simply disappear.

But we must not be too brusque in dismissing negative tests. EPAM-III, in fact, makes limited use of them for a reason that makes good psychological sense. Imagine that a subject is observing a large marmalade cat named "Cinnamon." If asked what he or she was looking at, the reply might be, "a cat," "a marmalade cat," or "Cinnamon." The latter answer depends, of course, on the subject recognizing the particular cat in question. Suppose the subject was asked, "What cat is that?" and replied, "None of the cats that I know." How can the subject *know* that the cat is not numbered among those known to him or her? We have seen in our description of its architecture that EPAM-III accounts for this response readily by providing the possibility of a negative, n.e.c. (not elsewhere classified), branch at every node. Thus, feature tests at the marmalade-cat node of the discrimination net could separate out various known cats on the basis of their specific features, but the n.e.c. node would be associated with the memory contents for the class of marmalade cats, and would be accessed by the discrimination process whenever the cat under observation was an unfamiliar marmalade cat (i.e., did not satisfy the test at that node for any branch). Hence, we cannot understand Barsalou and Bower's claim that EPAM "would produce the same action to unfamiliar stimuli . . . as to familiar stimuli" (p. 6).

Negative information is very important, therefore, in telling us that we cannot refine the discrimination we have made beyond some general class of objects, and provides the means for accessing directly the information we have about particular classes at any level of abstraction. Depending upon circumstances, we may recognize a "furry object," a "mammal," a "cat," a "marmalade cat," or "Cinnamon." Thus, negative properties serve the important, if limited, function of preventing us from continuing down the discrimination tree below the level of our information. They do not prevent all misidentifications, since the net does not contain comprehensive and conclusive tests for the objects in it; but they permit us to know that a particular stimulus, while belonging to a class that we recognize, is not a partic-

ular familiar member of that class (or, in other circumstances, that it does not belong to an already familiar subclass of that class). A grave difficulty with PANDEMONIUM-like systems is that they respond to the highest level of similarity detected; hence, it is not easy to see how they recognize that a stimulus is unfamiliar.

Barsalou and Bower also object to the assumption in EPAM that at the bottom of the net for each familiar stimulus there is stored a partial image of that stimulus, typically containing slightly more information than is contained in the tests that identify it. To those authors, this mechanism appears "unnatural and inefficient" (p. 5). But if they had attended to the task of *learning* discriminations, and especially to the task of learning to discriminate between two closely similar stimuli that are never observed simultaneously, they would see that this mechanism is both natural and efficient.

Consider the task of learning to distinguish two identical twins who are never seen together at the same time. Each time a misidentification is made, some additional feature or features of the twin who is present must be added to the discrimination net. But how is it known that the additional feature will be discriminative? Out of all features that might be noticed, only a few discriminate between the twins. If the image at the terminal node of the discrimination net contains information about additional features of the absent twin, then that information can be compared with the features of the twin who is present, and particular features selected for extending the net that will discriminate between them. If there are no such discriminating features in the image, then one or more features are added to it, and the same process is repeated when one of the twins next appears.

In any event, because of the obvious difference in use and structure of an "index" and an "encyclopedia," it would be unnatural to use a single representation for identifying a pattern and for storing knowledge about its properties. To the best of our knowledge, none of the current models of semantic memory make this identification.

SENSITIVITY TO DISCRIMINATIVENESS OF PROPERTIES

Barsalou and Bower claim that EPAM nets do not give preferential treatment to those properties of stimuli that are likely to discriminate among them. As "proof," they show a hypothetical EPAM net in which the tests are non-optimal. What they do not show (and cannot show, since it is not true) is that this is the kind of net that will emerge from EPAM's learning experiences. An EPAM net is not just any arbitrary discrimination net; it is a net that has grown and differentiated on the basis of learning experiences.

EPAM typically learns (expands the discrimination net) as the result of confusing an unknown stimulus with one previously learned. It recognizes its mistake because of the redundant information (mentioned in the

previous section) that is stored at the terminal node to which it has (erroneously) sorted the new stimulus. By comparing the novel stimulus with the incomplete image of the familiar stimulus, it discovers a feature that will discriminate between them, and inserts precisely *this* feature into the net. Hence, contrary to the claim of Barsalou and Bower, EPAM is exquisitely sensitive to learning just those features that will have great discriminative value. EPAM, like the subjects in Barsalou and Bower's (1980) demonstration experiment, would recall best (from the image) those symptoms of a disease that discriminated it (in the net) from other diseases, since these are just the features that would be stored first in the image. There is no guaranty, of course, (in EPAM or human) that the choice will be optimal.

EPAM also sometimes learns by encountering at some node a value for the feature it is testing that does not correspond to any previous branch at that node. It then creates a new branch for the new value. We have already observed that it is this possibility of encountering new values not corresponding to any given branch in the net that enables EPAM to know that it has failed to recognize the stimulus.

Barsalou and Bower also object that the noticing order EPAM uses limits its ability to select the features for discrimination that are most informative. There are several answers to this objection. First, the noticing order postulated in EPAM-I and EPAM-II is under strategic control, and can be altered by the subject under particular circumstances or as a result of experimental instructions (as we have already remarked in discussing the serial position curve). Second, there is a large amount of evidence that the noticing order used in EPAM-III does describe correctly the usual sequence of acts of attention to verbal stimuli. The serial position curve data provide strong evidence for it, as do the spelling errors predicted by the SPEL program, and the data cited by Brown and McNeill (1966) in their paper on the tip-of-the-tongue phenomena. Now the latter authors also show that the beginnings and ends of words contain more information per letter than the middles—which might account, at least partly, for the emergence of this particular sequential strategy. But as we have already shown, EPAM does not rely mainly on noticing order in choosing features for discrimination. This choice is mainly determined by comparing stimuli with the images with which they are confused.

Barsalou and Bower report an unsuccessful attempt to modify Hintzman's (1968) SAL net learning scheme to make it more discriminative. However, the mechanisms they introduced bear little resemblance to EPAM's use of the redundant stimulus image and, moreover, they evidently employed a binary net in their experiment. Therefore, neither this experiment and simulation, nor the authors' analysis tell us anything about the discriminativeness of the tests that EPAM incorporates in its nets.

SENSITIVITY TO MISSING OR INCORRECT PROPERTIES

It is claimed that EPAM nets are too sensitive to missing or incorrect information. The argument is supported, in the Barsalou and Bower paper, by reference to a hypothetical binary net that is wholly irrelevant to the multiply branching nets actually employed by EPAM-III. We have already explained how the n.e.c. branches at nodes in the EPAM-III nets are able to handle missing information. The new cat on the block is not, in fact, recognized erroneously as one of the familiar cats, but is simply recognized as a cat, or as a marmalade cat. With respects to incorrect properties, much more can be said.

Recognition Errors

First, when people identify features incorrectly, they frequently make recognition errors. Psychological research has made good use of confusion matrices, which are eloquent testimony to the human capacity for making mistaken identifications. A model of psychological processes that did not make errors of this kind would not be at all veridical.

EPAM-III, having multiple branches at each node and making little use of negative information, will confuse stimuli with other stimuli that are similar, and not with totally dissimilar ones. Barsalou and Bower's highly simplified and idiosyncratic example based on a tiny binary net is wholly misleading.

Second, there is no requirement in the EPAM theory that there be a single correct path of tests leading to each distinguishable object. The EPAM net is an index to the information stored in long-term memory, and that index can be highly redundant, with many paths leading to the same terminal node. This feature of EPAM nets has not been much studied, although it was an important aspect of the nets employed in the MAPP program to which we have previously referred. A plausible interpretation of the well-known phenomenon of overlearning is that it involves the creation of redundant access paths so that information already stored in memory is more readily recovered (Trabasso & Bower, 1968).

Persons and objects are viewed, on different occasions, from many different angles, and views may be partial and incomplete in many ways. The noticing mechanisms that control attention will pick up quite different features on these different occasions. Multiple paths—and a great variety of them—are therefore essential to the kind of recognition that we experience in everyday life. Nor is it likely that this multiplicity of viewpoints could be handled by varying the weights in a common set of features in a single PAN-DEMONIUM-like "demon." The features of a friend that are used to rec-

ognize him when he is at a distance, walking away, have almost nothing in common with the features used when he is sitting across the luncheon table. Recognizing him on both occasions requires two separate “demons”—that is to say, two independent recognition paths.

Role of LTM in Recognition

Also, multiple paths arise in another way. Recognition may be direct (direct access to an object's schema in semantic memory); but it may also be indirect (access by recognition of something associated with the object in semantic longterm memory). Thus, if the stimulus that is observed is a wheel attached to something, and the features of the wheel include pneumatic tires and hubcaps, that something may be (correctly or incorrectly) recognized as an automobile; because in semantic memory, the automobile wheel scheme is associated, as “part of,” with the automobile schema. In most circumstances, a person would make no distinction between a “direct” recognition of an automobile, whatever that might mean, and “indirect” recognition by association with the appropriate kind of wheel. In fact, most recognitions of complex objects probably take place through recognitions of associated parts.

In this way, the robustness of the discrimination net as a recognition device is greatly increased by the myriad of associational connections in longterm memory among familiar, recognizable items. Was that object recognized as a cat because it was discriminated along one of the “cat” paths, or because a cat's face having been recognized, that feature then associated to “cat” in semantic memory? Only very precise latency studies could detect the degree of indirectness in a recognition.

This interaction between the recognition mechanism and longterm memory provides an explanation for a phenomenon that is noted by Barsalou and Bower: that the recognition latencies of experts are frequently shorter than the latencies of novices. Because of the greater redundancy of paths in the experts' discrimination nets, we would expect them to accomplish a recognition directly, without resort to LTM associations, more frequently than novices. And in the case of indirect recognitions, the richer network of associations available to experts in LTM would imply association paths that would, on average, be shorter than the novices' paths.

Context Effects

In addition to the context effects just discussed, deriving from the redundancy of the EPAM net and from the interaction of the discrimination process with association links in LTM, the recursive structure of EPAM introduces

additional context effects that are of importance in such tasks as letter perception (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982).

When a stimulus is presented to EPAM, the attentional strategy determines what portion of the stimulus will be selected for recognition, and with what noticing order. For example, if the letters WINE are presented, preceded and followed by one or more spaces, EPAM will undertake to recognize the entire word. In the course of discriminating it, letters will be recognized, and this recognition will be achieved, in turn, by tests upon features. If a non-word is presented (e.g., WONE), the second letter may be misidentified (because of partial tests in the word net), or identification may be delayed, because only after recognition in the word net fails does the strategy shift to attending to individual letters. Hence, it appears that EPAM can provide at least a qualitative explanation of the superiority of recognition of letters when they appear in the context of a word. The phenomenon is not unlike the Stroop effect, where a highly overlearned strategy of discriminating stimuli for meaning interferes with recognition of the stimulus's (semantically irrelevant) color.

McClelland and Rumelhart (1982) explain the context effects in letter perception in a different way, using a parallel, PANDEMONIUM-like recognition scheme with spreading activation. Their scheme apparently ignores letter position within words, and appears to be rather sensitive to the parameter settings, but is otherwise impressive in fitting the empirical data. Since EPAM has not been tested empirically with the same data, we cannot give a comparative evaluation of the two hypotheses for these tasks.

MULTIPLE KNOWLEDGE DOMAINS

Barsalou and Bower claim (p. 9) that "the more knowledge domains in which an EPAM net discriminates, the more inefficient and psychologically implausible it becomes." This claim rests on two invalid bases, each vividly illustrated in their article by a figure (their Figures 3 and 4).

Their Figure 3, which "exhibits the problem of extended negative properties in multiple stimulus domains" (p. 10) is based, again, on the invalid assumption that EPAM-III requires binary nets. Remove that assumption and the alleged problem has disappeared from the redrawn figure.

Their Figure 4, which "exhibits the problem of exponential node growth in multiple stimulus domains" (p. 12) is based on a different misconception of how EPAM-III actually works. Suppose that a particular node of EPAM discriminates the first letters of words, while another node discriminates the second letters. Is it necessary to elaborate at each of these nodes the whole set of elementary features that allow EPAM to discriminate among letters, as Barsalou and Bower's Footnote 2 claims and as their Fig-

ure 4 depicts? Certainly not. EPAM is constructed to organize itself in recursive fashion. (See our earlier description of recursion in our account of the EPAM-III architecture.) Once it has learned to discriminate among the letters of the alphabet, it will have built a single discrimination net to accomplish that recognition. Whenever the identity of a letter is used as a test to discriminate some object (e.g., a word) of which it is a component, EPAM will use this unique letter discrimination net to identify the letter and to select the correct branch in the sorting node for the word. Thus, the four subnets shown in their Figure 4 simply do not exist in the EPAM system, and the horrid prospect of "exponential node growth" is neatly avoided. In the light of this, Footnote 2 of the Barsalou and Bower paper is simply incomprehensible.

One additional word needs to be said about "exponential node growth." The number of nodes in an EPAM net grows only *linearly* with the number of objects discriminated—not exponentially or geometrically or even in polynomial fashion. In the worst case, the binary net, the total number of nodes is twice the total number of objects discriminated. In an n -branch net, the total number of nodes is only $n/(n-1)$ times the number of objects. Of course, if the net is redundant, the number of nodes will be larger. If each object is accessed by k paths, the total number of nodes will be multiplied by k . Thus, there is absolutely no basis for the claim that an EPAM-like system will generate a net that is disproportionate to the size of the set of objects it is able to discriminate. Redundant PANDEMONIUM nets will of course also contain k times as many "demons" as objects, where k represents the number of different feature configurations that allow recognition of the same object.

SERIALITY

The relative roles of serial and parallel processes in human mental functioning has long been under debate. Very strong evidence can be given for the presence of both, and there is nothing in the EPAM system that is contradictory to that evidence.

EPAM and Parallelism

Let us recall that EPAM is not a model of the entire sensory-perceptual-memory system. It has always been explicitly assumed that it operates in conjunction with a sensory-perceptual "front end," and a semantic node-link longterm memory to which it gives access. The front end analyzes the stimulus into a set of features which it provides to EPAM. EPAM uses these features to index longterm semantic memory. Neither the front end nor the

semantic memory (except for the images at terminal nodes of the EPAM net) is part of the system that is modeled.

We do not believe that there is any contemporary dispute that the first stages of sensation and perception are essentially parallel in operation. The retina and the cochlea are most obviously parallel devices that extract a panoply of basic features from incoming stimuli. EPAM is a proposed model of what happens to those features once extracted, and how they are used for the recognition of objects.

On the other hand, there is also ample evidence that human attention is severely limited, so that only a very few tasks can occupy attention "at one time." It is not an unreasonable hypothesis, certainly one not contradicted by any substantial evidence, that only a single task, and a few chunks of information can occupy attention in a single "moment." Where several tasks are attended to over a period of seconds, this phenomenon can be accounted for by time sharing. But the EPAM theory could easily be adapted to the alternative assumption that a few tasks could be handled at once—as long as "few" was a modest number, and "at once" meant over some fractions of seconds.

With respect to the discrimination net itself, even larger concessions to parallelism require little modification of the system. In fact, in the case of a highly redundant EPAM net, alternative paths that were simultaneously stimulated could be followed concurrently, and it is very hard to think of tests that would distinguish such a net from the discrimination system of PANDEMONIUM or the "data flow" system proposed by John Anderson (1983, pp. 137–145).

Need for Seriality

But the discrimination net is only one component of the EPAM theory, and the assumption of seriality plays a more important role in connection with some of its other features—especially with respect to the anchor-point assumption that determines the order of learning of components of linear structures like words. Here, seriality is both essential and strongly supported by empirical data—as we have already seen.

Thus, the seriality in EPAM's noticing processes accounts for the empirical phenomena of the serial position curve, as well as for the fact, observed both from the tip-of-the-tongue phenomenon and from spelling errors, that words are discriminated largely on the basis of their first and last letters, and with little attention to those near the middle. Moreover, it explains these phenomena without any need for assumptions about varying "strengths" in the connections between particular features and the objects being recognized. In this way, it gets along with far fewer degrees of freedom than do systems that require a strength parameter for each path from feature to recognition demon.

Recognition Latencies

But let us return to the serial discrimination net and the latency data that are alleged to contradict its predictions. First, what estimate can we make of the time it might take EPAM to make a discrimination using n -branch tests? Suppose that $n = 8$ (there are eight branches, on average, at each node), and that each test takes 10 ms. Then seven tests (taking a total of 70 ms) will discriminate among two million stimuli ($8^7 = 2,097,152$). In a net 10 tests deep (100 ms), over 1 billion objects can be discriminated. There is no evidence of recognition speeds of experts faster than these, nor any evidence that experts can discriminate among billions of different things. Even a binary net can discriminate among a million stimuli in 200 ms.

From these circumstances it is clear that the growth in depth of experts' discrimination nets could add only a few tens of milliseconds to their recognition latencies. But, as we saw in the previous section, these increases would be much more than compensated by the greater frequencies with which experts could make direct recognitions, without recourse to LTM associations, as well as the shorter, faster paths that would be available in LTM in the cases of indirect recognitions. There are other reasons, too, why experts' responses will be faster. In tasks like those studied by Shiffrin & Schneider (1977; Schneider & Shiffrin, 1977), where automation leads to greater speed, responding requires more processes than just simple recognition, and there are many opportunities for experts to compile subprocesses that novices perform interpretively.

COMPARISON OF EPAM WITH ALTERNATIVES

Barsalou and Bower sum up all the alleged deficiencies of EPAM under the label of "test contingency." We have seen that, even ignoring the scantiness of the empirical evidence they bring to bear, the principal criticisms that Barsalou and Bower make of EPAM simply do not hold water. EPAM does not make the excessive use of negative properties that they claim, although it does use the n.e.c. branch at each node to good purpose. Contrary to their claim, EPAM is indeed sensitive to the discriminativeness of properties. EPAM's alleged "brittleness," through sensitivity to missing or incorrect properties, is remedied (as in humans) by generating a redundant discrimination net acquired by overlearning. EPAM's supposed inefficiency in the presence of multiple knowledge domains vanishes as soon as we substitute n -way branches for binary branches, and the alleged exponential explosion of nodes never occurs—node growth is only linear with discriminating power. Finally, EPAM's seriality, far from being a defect of the model, is supported by a large body of evidence on human cognitive performance.

Specifically, we have dealt with all eight sets of empirical evidence that Barsalou and Bower cite against EPAM. Two sets of citations (their pp. 4 and 5) show that human subjects (like EPAM-III!) do not make extensive use of negative properties to describe stimuli. Two (pp. 6 and 7) show that human subjects, like EPAM, learn features of stimuli that are relatively discriminative. One (p. 8) shows that, like EPAM, human subjects sometimes confuse stimuli that are similar. Three are arguments for parallelism. Of these, one (p. 11) refers to the speedup of recognition with practice. One (p. 11) cites the neurological evidence for parallelism of the *feature extraction system* (wholly consistent with EPAM). The third (p. 19) is the McClelland-Rumelhart model that we discussed explicitly in connection with context effects. If there are empirical anomalies that confront the EPAM theory, these pieces of evidence do not describe them.

This refutation of the arguments raised against EPAM-like systems as models of the recognition process and the indexing of memory, together with the impressive positive record of matching experimental results create a strong presumption that EPAM provides a good explanation of these phenomena. However, before we are fully satisfied, we should compare EPAM with other candidate theories of recognition processes. The two principal contenders are PANDEMONIUM-like models and data flow models.

EPAM and PANDEMONIUM

The PANDEMONIUM architecture for a recognition system, whose principal author was Oliver Selfridge, was conceived at approximately the same time as EPAM. As usually described, it is narrower in its range of functions than EPAM. It does not deal with the strategies of attention or noticing order, but deals impartially with all of the stimulus information that is presented to it.³ Second, it learns to make discriminations only among a pre-specified set of alternatives; it does not gradually elaborate the fineness of the distinctions it makes among objects. Hence it is not possible to use the PANDEMONIUM architecture, without substantial modification, to simulate human performance in verbal learning experiments and to compare that performance with EPAM's.

PANDEMONIUM has two major components: a set of feature extractors that take the stimulus as input and produce a list of features contained in it; and a set of "demons," one for each of the classes of stimuli

³ Barsalou and Bower claim (their Footnote 4, p. 21) that PANDEMONIUM's insensitivity to order is easily remedied, presumably by adding more demons. They do not address the severe problem of combinatorial explosion that arises if one must represent by separate Demons all permutations of order. Real-world stimuli have orderly structure, and one pays an enormous penalty for parallelism that ignores that structure.

PANDEMONIUM is capable of recognizing. Each demon matches the stimulus features against a weighted vector of features and produces a measure of "goodness of fit." The demon that fits the stimulus best is then recognized. The set of features is given in advance, although PANDEMONIUM could be augmented by a mechanism for generating new features.

PANDEMONIUM operates in completely parallel fashion. In network terms, each feature that is capable of being detected may be considered to be connected with each demon, the strength of connection corresponding to the weight of that feature in the definition of the demon. Learning—improvement in accuracy of discrimination among a given set of demons—can take place through alteration of the strengths of connections.

To allow PANDEMONIUM to learn *new* discriminations, it would have to be able to distinguish between selecting the *wrong* demon for a stimulus and selecting the *same* demon for two or more distinct classes of stimuli. It is not clear just how this information would be provided, but suppose that it were. Then two demons could be created to replace the overgeneralizing one, and could be separately "tuned" to make the desired discrimination. How well such a procedure would stimulate human discrimination learning we do not know, since to the best of our knowledge it hasn't been tried.

Notice that the system of McClelland and Rumelhart (1981) discussed earlier in this paper operates with a PANDEMONIUM-like mechanism, and suffers from the limitations we have just described. In its present form, it could not be used to simulate human verbal learning within the usual experimental paradigms.

The limits of PANDEMONIUM must be balanced against its attractive features. To many psychologists, its parallelism and its incorporation of "strengths" of connections seem "right." But for reasons we have given, it cannot, without substantial modification, be tested in the tasks with which EPAM has been so successful.

Data Flow Nets

The data flow net, developed most fully by Anderson (1983), has both EPAM-like and PANDEMONIUM-like features. It employs a sorting net which, however, is not wholly serial but permits several paths to be searched in parallel. It also associates weights with paths so that the search is not an all-or-none matter. It is not clear that the behavior of a data flow net, as described by Anderson, could be distinguished empirically from a redundant EPAM net, and in any event, data flow nets have not been tested in the rote learning environments that EPAM has simulated so faithfully. It would certainly not be difficult, and would probably prove informative, to create various kinds of hybrid architecture combining ideas from EPAM-like

discrimination nets and data flow nets and to study their behavior in such environments.

CONCLUSION

In the *1979 Annual Review of Psychology* (Simon, 1979, p. 378) one of us said, "The safest conclusion at the present time is that human LTM can probably be represented as a node-link memory with an EPAM like index, but that various alternatives are still open for the detailed structure and organization of that memory." After a careful examination of Barsalou's and Bower's criticisms of the EPAM architecture, we are prepared to repeat that judgement today. In fact the extensive experience that De Jong (1979a, 1979b) and Kolodner (1980, 1983a, 1983b, in press) have had, in the intervening years, with n -branch discrimination nets provides a considerable body of new evidence to support that earlier judgement.

We are sure that as knowledge of learning and memory processes continues to advance, our theories of these processes will undergo many changes. The original EPAM model and its modifications were generated in response to particular bodies of empirical data. Other investigators, focusing their attention on somewhat different phenomena, arrived at models like PANDEMONIUM or data-flow structures, and introduced mechanisms like association strengths and levels of activation. It is not uncommon, in the history of the sciences, for competing models of these kinds to reveal different, partial, aspects of reality, and to converge gradually toward more veridical theories that borrow heavily from a number of their predecessors. It seems most probable that the same thing will occur in this domain, and it would be foolish to treat any of these models as graven in stone.

What components of the EPAM theory seem at present to be most firmly established by evidence? First, there appears to be little controversy that the recognition-memory system is a tripartite structure consisting of a feature-extracting front end that operates in parallel fashion, a recognition or indexing system that uses these features to access longterm memory, and a semantic memory consisting of a myriad of nodes linked by directed associations. Although we have been focusing in this discussion (as the EPAM model does) upon the second of these three components, all of the theories we have referred to postulate, explicitly or implicitly, this tripartite structure.

When the main concentration of research on memory moved from verbal learning to free recall, as it did about a decade or two ago, the attention of psychologists shifted from the recognition process to association processes within the semantic longterm memory. Recognition mechanisms

like EPAM, and associative memory structures like HAM were sometimes described as though they were alternative theories for the same phenomena instead of complementary components of the tripartite structure outlined above. That was (and sometimes still remains) a misconception.

A second component of the EPAM theory that is strongly supported is the control structure—the set of mechanisms that determines what stimulus features will be noticed, and how noticing, in turn, will determine what discriminations are learned. That control structure, however, is substantially specialized to a particular class of stimulus materials—symbolic strings with spaces serving as string boundaries. When EPAM was used in the context of two-dimensional visual perception (chess perception), it had to be provided with an appropriate noticing mechanism (PERCEIVER) to control its attention to parts of the stimulus.

When we come to the details of the discrimination net itself, there is more room for debate—or preferably, for continuing empirical research. Discrimination nets and PANDEMONIUM-like structures both have attractive properties, and both appear to be able to account for at least the gross features of the recognition process. There is no reason why the choice between them has to be made in 1984—nor is it even clear whether we are faced with a choice or an assimilation. Rather, what seems to be called for is continuing vigorous development of both models (and others, as well), and equally vigorous search for experimental methods for obtaining finer data, especially quantitative data, about the recognition process. Meanwhile EPAM, or some variant of it, will continue to serve as a useful theory for summing up and explaining a very large proportion of the principal available for empirical findings about human recognition processes.

REFERENCES

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston & Sons.
- Atkinson, R. C., Bower, G. H., & Crothers, E. J. (1965). *An introduction to mathematical learning theory*. New York: Wiley.
- Barsalou, L. W., & Bower, G. H. (1980). *A priori determinants of a concept's highly accessible information*. Paper presented at the meeting of the American Psychological Association, Montreal.
- Barsalou, L. W., & Bower, G. H. (1984). Discrimination nets as psychological models. *Cognitive Science*, 8, 1-26.
- Brown, R., & McNeill, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior*, 5, 325-337.
- Bugelski, B. R. (1962). Presentation time, total time, mediation in paired-associate learning. *Journal of Experimental Psychology*, 63, 409-412.
- Chase, W. G., & Simon, H. A. (1973). The mind's eye in class. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press.

- De Groot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton.
- De Jong, G. (1979a). Prediction and substantiation: A new approach to natural language processing. *Cognitive Science*, 1979a, 3, 251-273.
- De Jong, G. (1979b). *Skimming stories in real time*. Unpublished doctoral dissertation, Yale University.
- Feigenbaum, E. A. (1959). *An information processing theory of verbal learning*. Santa Monica, CA: RAND Corporation.
- Feigenbaum, E. A. (1961). The simulation of verbal learning behavior. *Proceedings of the Western Joint Computer Conference*, 19, 121-132.
- Feigenbaum, E. A. (1967). Information processing and memory. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 4, 37-51.
- Feigenbaum, E. A., & Simon, H. A. (1961). Forgetting in an associative memory. *Proceedings of the ACM National Conference*, 16, 202-205.
- Feigenbaum, E. A., & Simon, H. A. (1962). A theory of the serial position effect. *British Journal of Psychology*, 53, 307-320.
- Fisher, R. P. (1981). Interaction between encoding distinctiveness and text conditions. *Journal of Experimental Psychology: Human Learning and Memory*, 1981, 7, 306-310.
- Gregg, L. W., & Simon, H. A. (1967). An information-processing explanation of one-trial and incremental learning. *Journal of Verbal Learning and Verbal Behavior*, 6, 780-787.
- Hintzman, D. L. (1968). Explorations with a discrimination net model for paired-associate learning. *Journal of Mathematical Psychology*, 1968, 5, 123-162.
- Hull, C. L. (1935). The influence of caffeine and other factors on certain phenomena of rote learning. *Journal of General Psychology*, 13, 249-273.
- Kolodner, J. L. (1980). *Retrieval and organizational strategies in conceptual memory: A computer model*. Unpublished doctoral dissertation, Yale University.
- Kolodner, J. L. (1983a). Maintaining memory organization in a long term dynamic memory. *Cognitive Science*, 7, 243-280.
- Kolodner, J. L. (1983b). Reconstructive memory: A computer model. *Cognitive Science*, 7, 281-328.
- Kolodner, J. L. (in press). *Retrieval and organizational strategies in conceptual memory: A computer model*. Hillsdale, NJ: Erlbaum.
- Krumhansl, C. L. (1982). Density versus feature weights as predictors of visual identifications: Comment on Appleman and Mayzner. *Journal of Experimental Psychology: General*, 111, 101-108.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos, & A. Musgrave (Eds.), *Criticism and the growth of knowledge*. Cambridge: Cambridge University Press.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive model of context effects in letter perception, Part I: An account of basic findings. *Psychological Review*, 88, 375-407.
- McCrary, J. W., & Hunter, W. S. (1953). Serial position curves in verbal learning. *Science*, 117, 131-134.
- Miller, G. A. (1956). The magical number seven, plus or minus two. *Psychological Review*, 63, 81-97.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Quillian, M. R. (1966). *Semantic theory*. Unpublished doctoral dissertation. Pittsburgh: Carnegie Institute of Technology.
- Quillian, M. R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing*. Cambridge, MA: MIT Press.
- Popper, K. R. (1959). *The logic of scientific discovery*. London: Hutchinson.
- Popper, K. R. (1965). *Conjectures and refutations: The growth of scientific knowledge*. New York: Harper & Row.

- Rock, I. (1957). The role of repetition in associative learning. *American Journal of Psychology*, 70, 186-193.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.
- Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception, Part II: The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, 89, 60-94.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review*, 84, 1-66.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 127-189.
- Simon, D. P., & Simon, H. A. (1973). Alternative uses of phonemic information in spelling. *Review of Educational Research*, 43, 115-137.
- Simon, H. A. (1974). How big is a chunk? *Science*, 183, 482-488.
- Simon, H. A. (1979). Information processing models of cognition. *Annual Review of Psychology*, 30, 363-396.
- Simon, H. A., & Barenfeld, M. (1969). Information-processing analysis of perceptual processes in problem solving. *Psychological Review*, 76, 473-483.
- Simon, H. A., & Feigenbaum, E. A. (1964). An information-processing theory of some effects of similarity, familiarization, and meaningfulness in verbal learning. *Journal of Verbal Learning and Verbal Behavior*, 3, 385-396.
- Simon, H. A., & Gilmarin, K. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, 5, 29-46.
- Trabasso, T. R., & Bower, G. H. (1968). *Attention in learning: Theory and research*. New York: Wiley.
- Winograd, E. (1981). Elaboration and distinctiveness in memory for faces. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 181-190.