

Simulations of Classification Learning using EPAM VI

Howard B. Richman and Herbert A. Simon*

Complex Information Processing
Working Paper #552

Department of Psychology
Carnegie Mellon University
Pittsburgh PA 15213

March 4, 2002

*Herbert A. Simon was involved with the formulation of the model reported here but died on February 9, 2001, before the simulations were conducted or this report was written.

“We have come gradually to the conclusion that what is most needed in the analysis of categorizing phenomena...is an adequate analytic description of the actual behavior that goes on when a person learns how to use defining cues as a basis for grouping the events of his environment.”

(Bruner, Goodnow, & Austin, 1956, p. 23)

Please send correspondence to: Howard B. Richman
RR 2 Box 117
Kittanning PA 16201
Email: howard@pahomeschoolers.com
Phone: 724-783-6512

Copyright by Howard B. Richman.
Copyright for some of the included
figures is held by the journals where
they first appeared.

Abstract

In the simulations reported here, EPAM VI explained classification learning experiments from Nosofsky, Gluck, Palmeri, McKinley, & Glauthier (1994), Martin & Caramazza (1980), Nosofsky, Palmeri, & McKinley (1994), Medin & Smith (1981) and Palmeri & Nosofsky (1995). The results support the idea that people form and test hypotheses about features and conjunctions of features and then memorize successful hypotheses (that they think will allow the number of unclassifiable items to be reduced) in a discrimination net. EPAM explained results earlier explained by Nosofsky, Palmeri, & McKinley's (1994) RULEX as well as results that RULEX would have difficulty explaining including: (1) the response latency times that result from the discrimination nets inferred by Martin & Caramazza for their subjects, (2) the longer-latency times under rules-plus-exceptions instructions of Medin & Smith's subjects, and (3) the old-new recognition ratings of Palmeri & Nosofsky's subjects.

When Rosch and Mervis (1975) demonstrated that people could learn family-resemblance classification structures, they turned the field of categorization learning on its head. Previous researchers, including such pioneers as Bruner, Goodnow, and Austin (1956), had assumed that a category needed to have defining cues and were busy exploring how people learned such cues. But after Rosch and Mervis, the new models that were developed, including the very successful context model of Medin and Schaffer (1978), used probabilistic response rules, not defining cues, to classify family resemblance stimuli.

Such models reigned, unchallenged, until Nosofsky, Palmeri, and McKinley (1994) introduced RULEX, a model of classification learning that could almost do it all. Its basic assumption was that people learn a defining cue (called a rule) and exceptions to that rule. Although rooted in the early criterial-characteristic categorization models of the field, it could beat the new probabilistic-response-rule models at the new family-resemblance tasks. It was clearly a great advance. But RULEX ran into two road blocks.

The first road block countered one of RULEX's premises, the supposition that people solve family-resemblance categories by first finding a rule and then learning exceptions to the rule. Even before the authors had published the first paper about RULEX, Medin and Smith (1981) had found that people do not appear to use rules-plus-exceptions instructions as their standard strategy when they do a classification learning task.

The second road block was noticed by authors of the RULEX model themselves when Palmeri and Nosofsky (1995) found that they couldn't use RULEX to predict recognition data without merging it with Medin and Schaeffer's (1978) probabilistic-response-rule context model.

In this paper we present EPAM VI, another deterministic model which also reaches back into the past for solutions to problems of the present. We will show that, when we add a RULEX-like strategy to EPAM VI, it is able to simulate what RULEX was earlier able to simulate. But because EPAM stores its learning in a discrimination net, not as rules and exceptions, we will show that it is able to get around both of RULEX's road blocks. Thus we will build on the progress made by RULEX, but go a bit further down the road.

EPAM VI shares many characteristics with Hunt, Marin, and Stone's (1966) CLS model, the first discrimination net model of classification learning (Richman, 1991). Both EPAM and CLS hold the outcome of their learning in a discrimination net that is built over a simulated-subject's run of a classification learning experiment. Both use a strategy that chooses the test at each node based upon similarities within a category.

Both can utilize either simple or compound tests at a node. Perhaps the only differences are task and strategy. CLS used a strategy that worked well within the non-incremental classification learning tasks that it simulated. EPAM VI uses a strategy that is designed to work in the typical incremental classification learning design in which:

1. The categories are *artificial* – in other words, classifications are arbitrarily determined by the experimenter.
2. The learning is *supervised* – in other words, the subject is given feedback after each classification try.
3. The stimuli are presented *incrementally* – in other words, the subject only sees one stimulus at a time.
4. The categorization is *binary* – in other words, there were only two category choices, usually categories “A” and “B,” and the subject’s task is to classify each stimulus as either Category A or Category B.

So that EPAM could utilize a strategy for choosing a test at a test node, we had to break with earlier versions of EPAM which automatically chose tests through a process of comparing the stimulus presented with a stored image of a previous stimulus. In EPAM VI, the choice of test is subject to strategic control. In the other experimental paradigms that we have simulated using EPAM VI, the stimuli were ordered lists: either lists of letters (i.e. nonsense syllables), lists of syllables or words, or lists of digits. In those experiments, we gave EPAM the anchor-point strategy (Feigenbaum & Simon, 1962) of studying such stimuli in an ordered fashion beginning either at one end of the list and working backwards or starting at both ends of the list and working inwards. In the classification learning experiments simulated here, we gave EPAM a hypothesis-formation strategy.

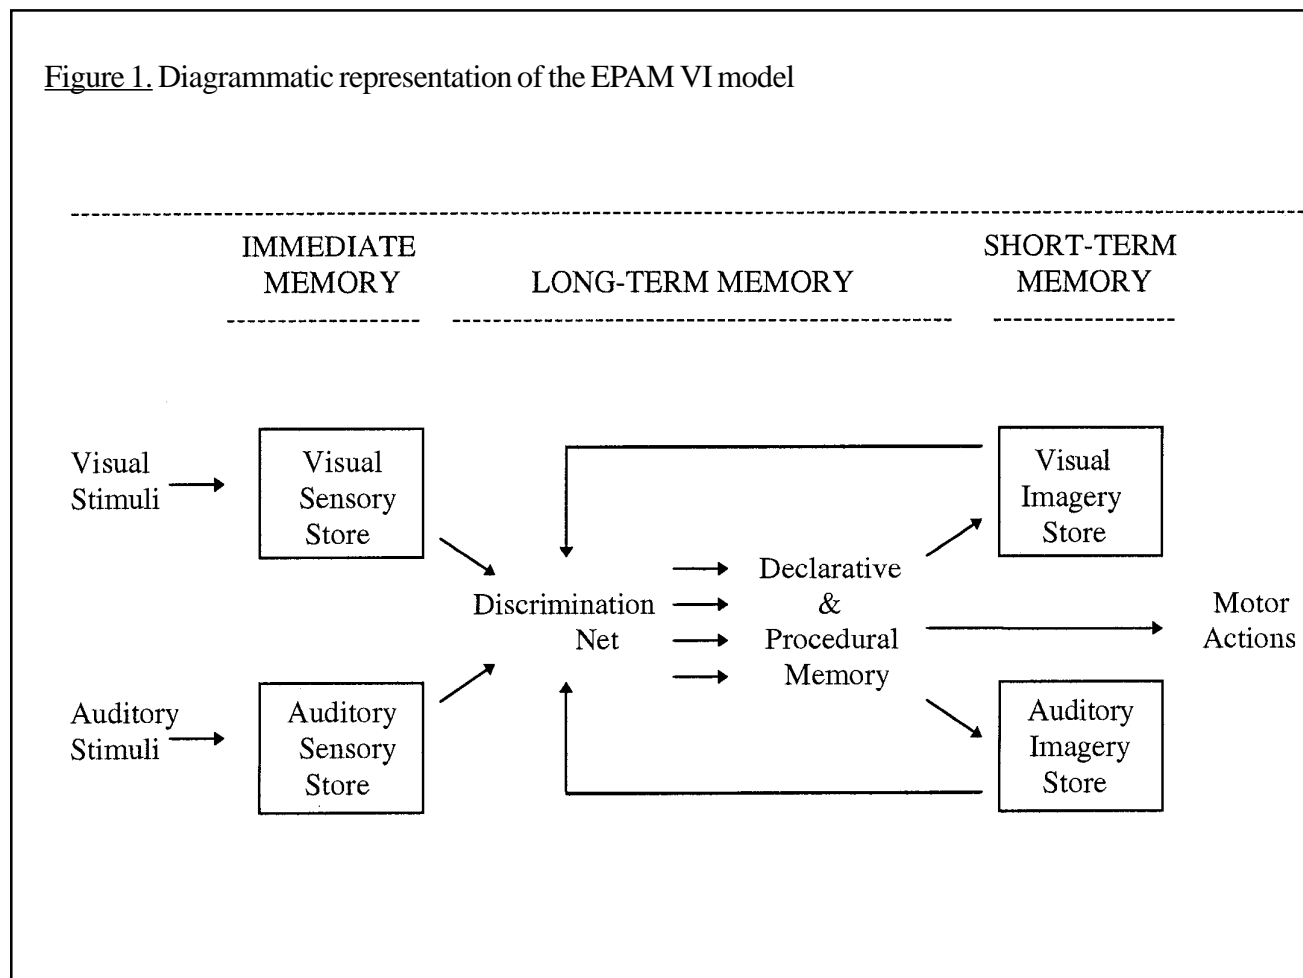
The EPAM VI Model

EPAM VI is the sixth revision of Elementary Perceiver And Memorizer, a model first developed by Feigenbaum (1959). Although EPAM VI models immediate memory, short-term memory, and long-term memory, its core is a part of long-term memory called the discrimination net which accesses the rest of long-term memory much as the index of an encyclopedia accesses an encyclopedia’s content.

Programming of the sixth version of EPAM was just completed in February, 2002, and our reports

about its simulations in a wide variety of experimental paradigms (including paired-associate, serial anticipation, fan-effect, distractor task, and auditory-span paradigms) are now in the process of being submitted for publication. The new revision combines EPAM III (Feigenbaum & Simon, 1984) and SAL III (Hintzman, 1968), the two leading models of the paired-associate paradigm, and continues to simulate the expert memory data earlier simulated by EPAM IV (Richman, Staszewski, & Simon, 1995) and the context effects in letter perceptions data earlier simulated by EPAM IIIA (Richman & Simon, 1989). EPAM VI thus provides an internally-consistent explanation of a wide variety of phenomena. Full information about the EPAM model, including the actual code (written in Allegro Common Lisp) can be found on the Internet at: www.pahomeschoolers.com/epam.

Figure 1 shows a diagrammatic representation of the EPAM VI model. The arrows in the figure show the flow of information. There are immediate memory stores and short term memory stores in each modality. The visual sensory memory corresponds to Sperling's (1960, 1963) icon, and the visual and auditory imagery



stores correspond to Baddeley's (1983, 1986, 1990) visuo-spatial sketch pad and articulatory loop). There is also a long-term memory consisting of a discrimination net, and Anderson's (1983, 1993) declarative and procedural memory. There is no clear demarcation between the discrimination net and declarative memory since discrimination net nodes are themselves part of declarative memory.

Stimuli enter the system through the sensory stores where they are sorted by the discrimination net to appropriate entries in declarative memory. Procedures stored in procedural memory act upon the outside world through motor actions, or upon the visual imagery store through visualizations, or upon the auditory imagery store through subvocalizations, or upon long-term memory itself through memorization and forming new associations.

Just as external information from the sensory stores can be sorted through the discrimination net, so can internal information from the imagery stores. For example, during subvocal rehearsal, information that is in the auditory imagery store is sorted through the discrimination net and then placed back into the auditory imagery store.

Discrimination Net

The core idea of EPAM is that perception is organized in a way akin to the 20 questions game. At each node of a discrimination net a question is asked about the to-be-perceived object. In this way, the object is sorted from one node to another down the net until a node that can function as a symbol for that object is located. New memories can be added by growing new nodes that can be accessed if the proper questions are answered. The first two versions of EPAM used binary ("yes" or "no") questions at each node. Later versions used n-ary questions (i.e. questions that can have any number of answers).

The top node of EPAM's discrimination net has a basic category test for "is-a." It, in effect, ties various EPAM subnets together. Every different type of object has its own subnet. For example, EPAM VI's discrimination net has a subnet for written words (is-a word) a subnet for spoken letter names (is-a letter), and a subnet for spoken number names (is-a number). In the simulations discussed in this paper, EPAM creates and develops new subnets during each experimental run.

Declarative Memory

But the discrimination net is not all that there is to long-term memory, it is just the index. Just as each word in the encyclopedia index may point to several entries in the encyclopedia, each node in a discrimination

net may point to several entries (i.e. nodes) in declarative memory. Our simulations of the fan-effect depend upon the same representation of declarative memory that Anderson (1974, 1983) used when he simulated fan-effect experiments using the ACT models.

In the simulations reported here, EPAM, creates entries in declarative memory each of which unites a discrimination net node with its category. For example, if we have created a discrimination net node that conjoins white & large, we might create an episode in declarative memory which joins that node with the Category A (i.e. containing the information that “white & large” belongs to Category A in a particular experiment). This entry would be indexed both by the discrimination net node for Category A and also by the discrimination net node for white & large. The strategy utilized by EPAM does not make use of this link from the category to the stimulus, but it could do so if asked to describe what members of Category A look like.

Procedural Memory

EPAM does not currently include a realistic model of procedural memory. Currently, its procedural memory consists of a set of algorithms that have been programmed into it by the researcher. A future combination of EPAM with ACT-R could provide EPAM with the realistic procedural memory simulated so well by the ACT-R model (Anderson, 1993).

Short-Term Memory

EPAM’s short term memory includes sensory stores and imagery stores in both the auditory and visual modalities. In the classification learning experiments simulated in this paper, we do not know where people are actually storing short-term memory information about the stimuli being studied. We guess that subjects are using the visual imagery store since the stimuli are pictorial and so we use the visual imagery store to hold either an hypothesis or similarities from one visual stimulus and its categorization to the next.

The hypothesis. Our model returns to the seminal ideas of Bruner, Goodnow and Austin (1956) that people learn concepts by forming and testing hypotheses. When present, the hypothesis, which EPAM holds in its visual imagery store,¹ consists of up to five items of information: (1) information necessary to sort to the discrimination net node in question, (2 & 3) up to two features of a stimulus, (4) the category hypothesized for those features, and (5) the confidence level of this categorization. An hypothesis is formed by comparing two stimuli, usually the current stimulus and the previously-presented stimulus. If the two have more than two features in common, the features used are chosen randomly using the relative weights of the features (which are

free parameters of the system) to determine the likelihood that one feature will be chosen before another.²

Similarities between consecutive stimuli. When the visual imagery store does not hold an hypothesis it holds similarities and whether the categories were identical from one stimulus to the next. Such information allows EPAM to create a new hypothesis. For example, if the previous stimulus was a large white triangle classified as Category A and was followed by a large white square also categorized as Category A, the short-term store would hold enough information to note that both stimuli were large and white and were categorized the same.³

Strategy

EPAM's strategy for doing these categorization tasks has two components: (1) a responding strategy, and (2) a learning strategy. The two work closely together. The learning strategy creates and tests hypotheses and adds information to EPAM's discrimination net. The responding strategy uses the hypothesis and the discrimination net to find the category.

The learning strategy that we use in this simulation grew out of our analysis of subject strategies in a replication that we conducted of Medin and Smith's (1981) Brunswik faces experiment (Gluck, Staszewski, Richman, Simon & Delahanty, 2001). We spent many hours analyzing the individual subject protocols and other individual subject results and were able to infer discrimination nets that most subjects appeared to form. From protocols that compared the present stimulus to the previous we could tell that at least 14 of 18 subjects in the verbal-protocol conditions at least once used a comparison to the last stimulus presented.

Some of the subjects who formed conjunctive tests appeared to notice pairs of features and their opposites. For example, some subjects who formed conjunctive tests for eye-height and mouth height spoke of long faces (high eyes and low mouth) and short faces (low eyes and high mouth) while others spoke of high faces (high eyes and high mouth) and low faces (low eyes and low mouth).

Although there was little evidence of overall strategy, we suspected, based upon our analysis of problem solving in other domains (Newell and Simon, 1972), that subjects were using a version of means-ends-analysis trying to find tests that would reduce the distance between the current state where they could not classify any stimuli successfully to a goal state in which they could classify all of the stimuli successfully. Given the family-resemblance category structure (the 5-4 task discussed later in this paper), they could not find any single feature (such as high eyes) which would successfully classify all of the stimuli having that feature. How-

ever when they checked for conjunctions of two features (such as long faces) they might find that all of the long-faced stimuli were members of Category A. There were still other stimuli left to classify, but at least they could then classify all long-faced stimuli correctly and they were satisfied that they were making progress toward a problem solution.

Based upon these observations and assumptions we programmed a strategy that was designed to incorporate the following two heuristics: (1) Look for a feature or a conjunction of two features that, when present, correctly classify some unclassified stimuli without misclassifying any of the yet unclassified stimuli, and (2) After a feature or pair of features is added to the net, check out the opposite value of that feature or pair.

In order to fit the human data closely, we had to add other heuristics as well. The strategy that we settled upon has two overall components, a responding strategy and a learning strategy, each of which has several alternative steps. The steps of these components and their interactions will be illustrated by detailed examples when we discuss EPAM's simulations of the first two Shepard, Hovland and Jenkins tasks.

Responding strategy. EPAM's simulated subject's responding strategy has four steps. The first to apply is executed and then the routine exits:

1. If the stimulus matches the hypothesis, EPAM makes the prediction stored with the hypothesis.
2. If the stimulus sorts to the same node as the hypothesis, but does not match the hypothesis, then EPAM predicts the category opposite to that of the hypothesis.
3. If there is no hypothesis or the stimulus does not sort to the same node as the hypothesis then try to make a prediction from long term memory using the discrimination net.
4. If all of the above fail, EPAM guesses the category randomly.

Learning Strategy. This is the algorithm that simulated subjects utilize after receiving the feedback of the correct response. The first step to apply is executed and then the routine exits:

1. If there is no subnet for this type of stimulus, then create one. This is a housekeeping job that is done by the learning routine.
2. If the stimulus matched the hypothesis and the response made by the simulated subject was incorrect, then delete the hypothesis from the visual imagery store.
3. If the stimulus matched the hypothesis and the response made by the simulated subject was correct, then either upgrade the confidence in the hypothesis (i.e. "possible" becomes "likely," "likely" be-

comes “very likely”) or, if the confidence was already “very likely” add the hypothesis to the discrimination net and at the same time create a node in the discrimination net for the opposite of the hypothesis. Once the hypothesis has been added to the discrimination net, it is deleted from the visual imagery store.

4. If the prediction made by the discrimination net led to an error and the node to which the stimulus sorts is not already marked as confused, then mark the node to which the stimulus sorts as confused.

5. If the following three conditions apply, then try to create an hypothesis: (1) there is not already an hypothesis, (2) the stimulus either sorts to a node that has been marked as “confused” or a node which has no associated category, and (3) both the current stimulus and the previous one were members of the same category. An hypothesis can be created if the stimulus and the previous one share at least one feature in common which is not already used by the stimulus to sort to its discrimination net node. If the stimulus and the previous one share no such features in common, then associate the node to which the stimulus sorts with the category of the stimulus.

6. This rule only applies to those simulations in which the subject has been instructed to study the exceptions to a rule. It is just like step 5 except that if the current stimulus and the previous stimulus are not members of the same category the hypothesis will be created by comparing the exception to itself instead of by comparing the exception (or non-exception) to the previous stimulus. This will allow hypotheses to be made about exceptions when they are not presented consecutively.

7. If the hypothesis is the same category as the stimulus but includes features that have a different value from the stimulus, then do one of the following if it applies: (1) if there were two features on the hypothesis and only one differs from the stimulus, then eliminate the feature that differs from the hypothesis, (2) if there was one feature on the hypothesis and there was no category associated with the node to which the stimulus sorts then associate the node with the category that is on the hypothesis and delete the hypothesis.

8. If none of the above apply, do nothing.

Parameters

The EPAM model has many parameters that do not change from simulation to simulation but only a very few parameters that can be adjusted to fit the data of a particular experiment. In this way it differs from models of classification learning such as the context model (Medin & Schaeffer, 1978) and RULEX (Nosofsky, Palmeri & McKinley, 1994) which have many parameters that can be adjusted from experiment to experiment

in order to fit the data.

Among those parameters that do not change are timing parameters which determine how long each act takes. The main timing parameter that applies in this experiment is the estimate that an act requiring an eye movement in a visuo-spatial domain requires 250 ms. This has been found true in reading (Just & Carpenter, 1987) and chess (De Groot & Gobet, 1996; Gobet & Simon, 2000). In the simulations reported here, eye movements caused by tests in the discrimination net largely determine response latency times. When a test in the discrimination net does not require an eye-movement, we estimate that it only takes 10 ms, not 250 ms.⁴ The other factor that affects response latency times is the time required to guess (step 4 of the responding strategy) which we have estimated, as in Gobet et al. (1997), to be 1000 ms.

Although EPAM also requires time to match to an hypothesis, that time does not enter into the latencies reported in this experiment because we have assumed that when subjects are instructed to respond as quickly as possible while maintaining accuracy, they discard the hypothesis perhaps because they know that without feedback they can no longer confirm or reject the hypothesis and that without the hypothesis they will be able to respond more quickly. Thus just the discrimination net is used for responding during the part of the experiment when latency times are collected.

EPAM also includes a few free parameters that can be adjusted from simulation to simulation. These include:

1. A *speed-of-learning* parameter which determines the percentage of time that the system engages in learning, given the opportunity. When less than 100, this parameter slows EPAM's learning. The idea behind this parameter is that some stimuli might be harder to hold in short-term memory than others and that some subjects may concentrate less well during an experiment than others.

2. Parameters that determine the weighting of each attribute. These parameters determine the noticing order that attributes are tried when forming an hypothesis. This noticing order is resampled every time that an hypothesis is created. If eye-height has a weight that is 16 times as high as the weight of nose-length, then 16 times out of 17, eye-height will precede nose-length in the noticing order that attributes are tried when choosing an hypothesis. In simulations of those experiments where physical dimensions are randomly assigned to logical dimensions these weights are all set equal to each other so that they no longer function as free parameters.

In this paper we will discuss the results that occur when EPAM VI is used to explain five classification experiments. First we will give examples of EPAM's strategies in action as we use EPAM to simulate Nosofsky, Gluck, Palmeri, McKinley & Glauthier's (1994) replication of Shepard, Hovland and Jenkins' (1961) classification learning tasks. Then we will compare EPAM's discrimination nets with the discrimination nets that were inferred by Martin and Caramazza (1980) from their subjects' post-experiment reports and latency times. Finally, we will simulate three experiments using Medin and Schaeffer's (1978) 5-4 task which reported: (1) results at both the group level and the individual level (Nosofsky, Palmeri & McKinley, 1994); (2) the effects of different strategy instructions (Medin & Smith, 1981); and (3) old-new recognition data. (Palmeri & Nosofsky, 1995).

Simulation of Shepard, Hovland and Jenkins Tasks


Shepard et. al.'s (1961) classification learning experiment compared six classification learning tasks using stimuli that had just three binary valued attributes. Thus the entire stimulus pool might consist of eight stimuli such as these:

1. ((is-a shj) (color black) (shape square) (size large))
2. ((is-a shj) (color black) (shape square) (size small))
3. ((is-a shj) (color white) (shape square) (size large))
4. ((is-a shj) (color white) (shape square) (size small))
5. ((is-a shj) (color black) (shape triangle) (size large))
6. ((is-a shj) (color black) (shape triangle) (size small))
7. ((is-a shj) (color white) (shape triangle) (size large))
8. ((is-a shj) (color white) (shape triangle) (size small))⁵

In the listing above, we have numbered the stimuli using the same numbering system that appears in Figure 2. Here are brief descriptions of the members of Category A in each task. The numbers in parentheses refer to the stimulus numbers in the figure.⁶

- I. Square (1, 2, 3, and 4).
- II. Black & square (1 and 2) or white & triangle (7 and 8).
- III. Square & large (1 and 3) or black & small (2 and 6).
- IV. Black & square (1 and 2), square & large (1 and 3), or black & large (1 and 5).

Figure 2. Top: The six types of categorization problems tested by Shepard, Hovland, and Jenkins (1961). The eight stimuli are denoted by the corners of the cubes. Assignments to categories are denoted by the ovals or rectangles that enclose the stimulus numbers. Bottom: Illustrative example in which the stimuli vary along the dimensions of shape (square vs. Triangle), color (black vs. White), and size (large vs. Small). [Note. This is Figure 1 from Nosofsky, Gluck et al., 1994]



- V. Black & square (1 and 2), square & large (1 and 3), or white & triangle & small (5).
- VI. Black & square & large (1), white & square & small (4), black & triangle & small (6), or

white & triangle & large (7).

We will discuss EPAM's data for each task separately, and then return to an overall discussion of how EPAM does each task. EPAM's simulations of all six Shepard, Hovland, and Jenkins tasks were run with 1000 simulated subjects, speed-of-learning-parameter set at 100 (the fastest speed), and weights for the three dimensions equalized because physical dimensions were randomly assigned to logical dimensions in the experiment simulated.

Task I

In Shepard, Hovland and Jenkins (1961) first task stimuli with a particular value of an attribute would be classified as Category A, and all of the stimuli with the other value for that attribute would be classified as Category B. For example, all of the stimuli whose color was "black" (stimuli 5, 6, 7, and 8) might be classified as Category A and those whose color was "white" (stimuli 1, 2, 3, and 4) as Category B.


EPAM's learning algorithm quickly finds the criterial feature. Here, for example, are the highlights of a typical run in which the criterial feature was "color," with "black" indicating Category A and "white" indicating Category B:

The first stimulus presented is a large white square ((is-a shj) (color white) (shape square)(size large))⁷ which is a member of Category B. There is no hypothesis and no association to a category in the discrimination net, so EPAM randomly guesses Category A, which turns out to be incorrect. The first step of EPAM's learning strategy kicks in, and EPAM creates a subnet for stimuli that are of the type "shj" which stands for "Shepard-Hovland-Jenkins." This will be the subnet of the discrimination net where learning will take place during this experimental run. The current discrimination net is shown in Figure 3.

The next stimulus presented is a large white triangle. EPAM again guesses wrong with Category B. Then, following step 5 of the learning strategy EPAM creates an hypothesis by comparing the current stimulus with the previous one using information that had been stored in the visual imagery store. The hypothesis that is created and stored in the visual imagery store is that a stimulus that is large and white predicts category B. This hypothesis is given an initial hypothesis rating of "possible."

The next two stimuli are both the same, two small black squares. These stimuli sort to the same node as the hypothesis (both sort to the top node of the subnet), so, due to step 2 of the responding strategy, Category A, the opposite category as the hypothesis is the response.

Figure 3. Discrimination net after subnet for stimuli that are type “shj” has been created.



The next stimulus to be presented is a small white square. It does not match the hypothesis since the size in the hypothesis is large. As a result it is misclassified (step 2 of the responding strategy) as a member of Category A, step 7 of the learning strategy kicks in causing the hypothesis to be narrowed to the one-feature hypothesis white predicts Category B. The hypothesis’s confidence level is still “possible”.

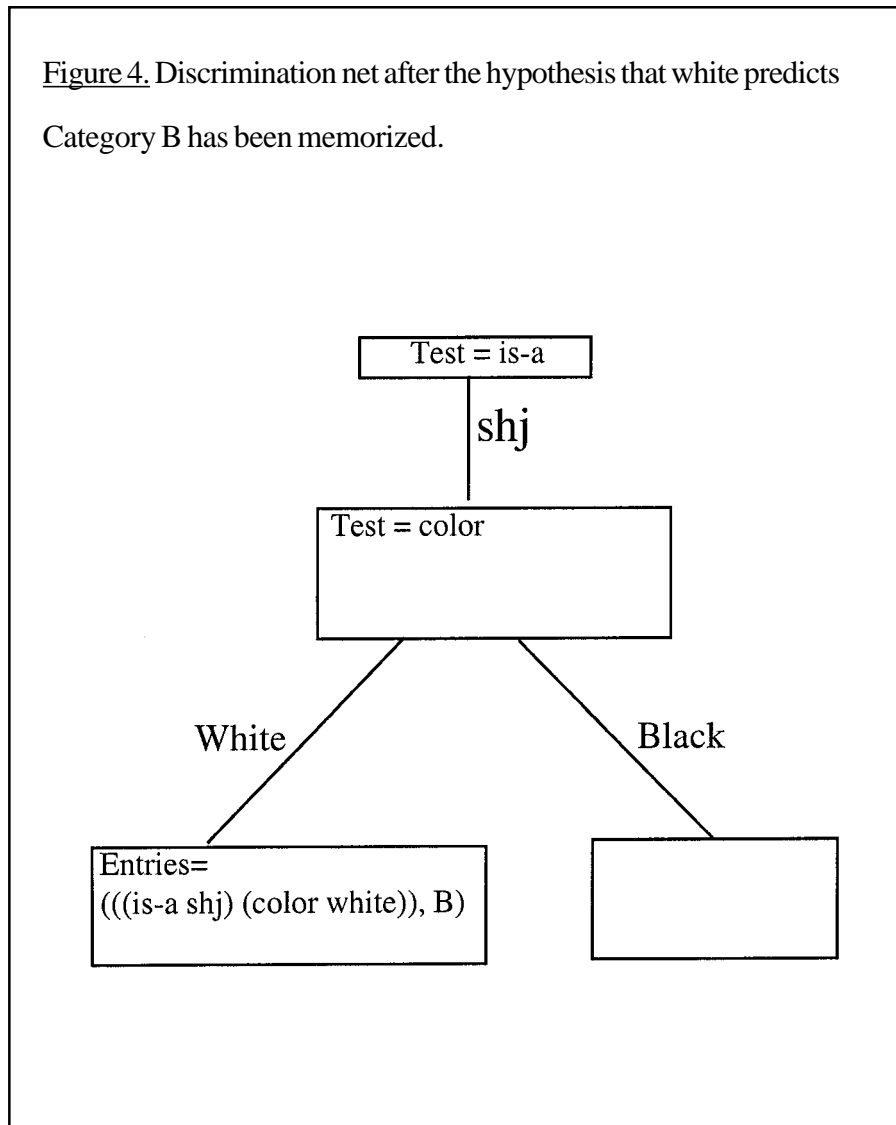
The next stimulus presented is a small white triangle. Since this stimulus matches the hypothesis EPAM responds with Category B which is correct, and so step 3 of the learning strategy raises the confidence level of the hypothesis from “possible” to “likely.” Similarly, the same step causes the next white stimulus to raise the confidence level to “very-likely.” When

the next white stimulus is presented, step 3 of the learning strategy causes the hypothesis to be added to the discrimination net and the opposite of its stimulus to be added to the discrimination net. The result is the discrimination net pictured in Figure 4.

When the next several stimuli are presented, EPAM correctly identifies white stimuli as Category B, but guesses whenever a black stimulus is presented. EPAM next engages in learning when a large black triangle is followed by a small black square. Both are members of Category A and so following step 5 of the learning process, EPAM tries to create an hypotheses. Since the two stimuli share no features in common (other than information that is already in the net for the stimulus) step 5 associates the node for “black” with Category A creating the discrimination net pictured in Figure 5. From this point on, EPAM responds correctly to every stimulus that is presented and engages in no further learning.

In their replication of Shepard, Hovland, and Jenkins (1961) experiment, Nosofsky, Gluck et al. (1994) reported the time course of learning for each of the six types of problems. The same data was also reported for RULEX. A comparison between the data for human subjects, RULEX (Nosofsky, Palmeri, and McKinley, 1994), and EPAM for Type 1 problems is shown in Table 1. Both models, like people, learn the concept in the first few blocks (each stimulus is presented twice during each block). Although EPAM makes

Figure 4. Discrimination net after the hypothesis that white predicts Category B has been memorized.



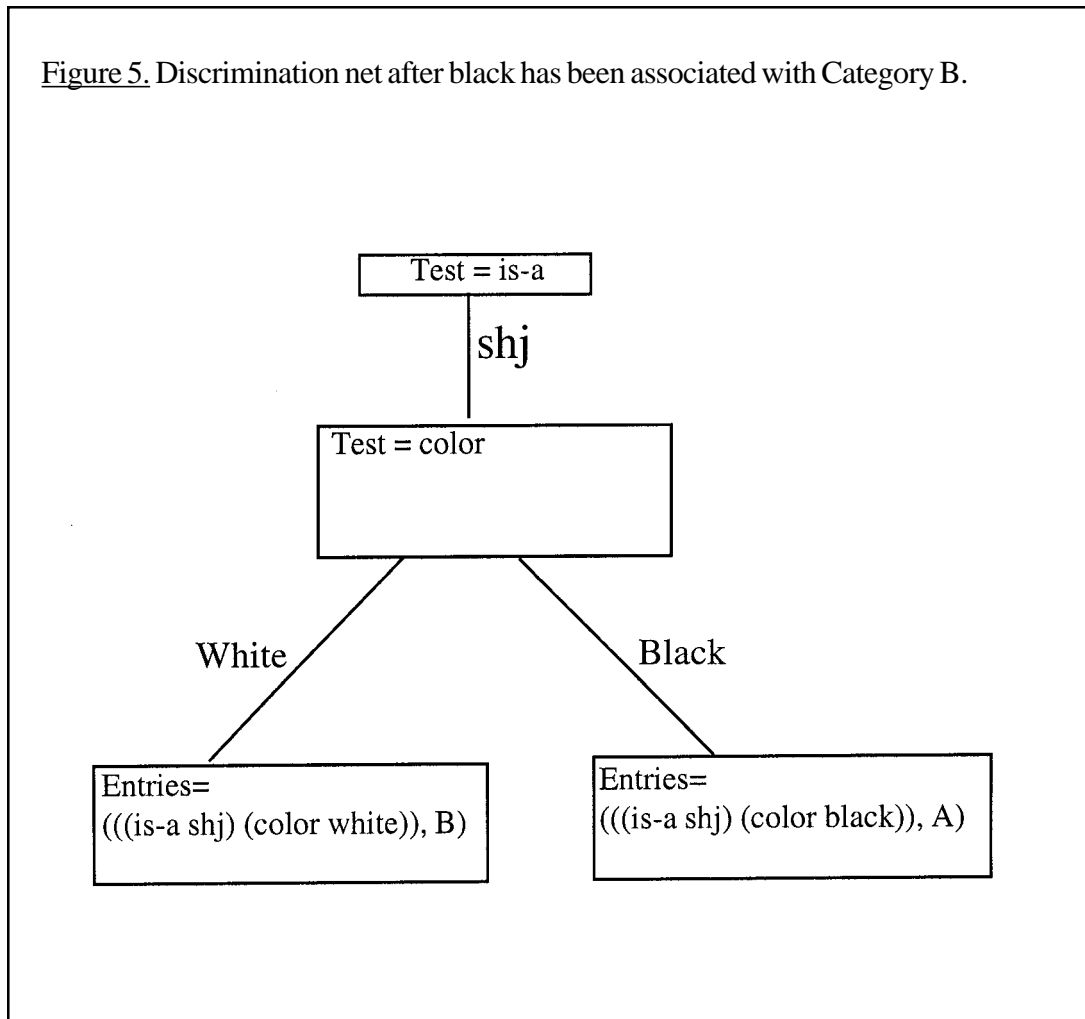
a few more errors than people in both the first and second blocks, it shows a perfect match with people after that.

The results shown in Table 1 have proven to be difficult results for some of the other models of classification learning. For example, Nosofsky, Gluck et al. (1994) found that two of the four models that they tested show a much more gradual decline in error rates on this task than do people.

Both EPAM and RULEX show the rapid decline in error rates over the first two blocks because they both use strategies that rapidly discover defining features. Once found, EPAM stores its solution as the top node of a subnet of a discrimination net, while RULEX stores its solution as a rule.

Task II

Figure 5. Discrimination net after black has been associated with Category B.



In Shepard, Hovland, and Jenkin’s (1961) second task stimuli that shared either two values or their two opposites were classified as Category A while all others were classified as Category B. EPAM’s learning algorithm first finds a conjunction of two features that correctly classify two of the stimuli without misclassifying any of the others. Then it checks out the opposite values of the test and discovers that it can classify the other two stimuli.

Here is a sample EPAM run for a set of stimuli in which either white & square or black & triangle (stimuli 3, 4, 5, and 6) would be classified as Category A while white & triangle or black & square would be Category B (stimuli 1, 2, 7, and 8). EPAM begins, as it did in the Task I example with step 1 of the learning strategy creating a subnet for “shj” resulting in the discrimination net pictured in Figure 3.

When two members of Category B (a small black square followed by a small white triangle) are presented, step 5 of the learning strategy creates an hypothesis with the confidence “possible” that small

predicts Category B. A few presentations later, a small white triangle is again presented and is correctly categorized by the hypothesis as a member of Category B. Step 3 of the learning strategy then raises the confidence in the hypothesis from “possible” to “likely.” However when a small white square is presented, and is incorrectly categorized by the hypothesis as a member of Category B, step 2 of the learning strategy deletes the hypothesis.

A few presentations later, a large white square follows a small white square and step 5 of the learning strategy creates the “possible” hypothesis that white square predicts Category A. Confidence in this hypothesis is raised to “likely” a few presentations later (step 3 of the learning strategy) when a small white square is presented and to “very likely” when a small white

square is presented again. Then when a large white square is presented, the hypothesis is added to the discrimination net and the opposite to the stimulus part of the hypothesis is also added creating the net pictured in Figure 6. From this point on, EPAM will always respond correctly to white squares.

A few presentations later two large white triangles are presented, one right after another. Following step 5 of the learning strategy, EPAM creates the “possible” hypothesis that large triangles are members of Category B. When a small white triangle member of Category B is next presented, EPAM narrows the hypothesis (step 7 of the learning strategy) to one that predicts that triangles are members of Category B. Confidence in this hypothesis is increased to “likely” when a small white triangle is again presented, but when a large black square is presented, step 7 of the learning process causes the hypothesis to be discarded and associates the top node of the “shj” subnet with Category B creating the net pictured in Figure 7. At this point,


Table 1

Average Error Proportions for Type I problems, in Blocks 1-6 of Learning

Block	People	RULEX	EPAM
1	0.21	0.19	0.24
2	0.03	0.01	0.06
3	0.00	0.01	0.00
4	0.00	0.01	0.00
5	0.00	0.00	0.00
6	0.00	0.00	0.00

Note. Human data is from Nosofsky, Gluck, et. al. (1994). RULEX data is estimated from the graph in Figure 5 of Nosofsky, Palmeri, and McKinley (1994).

Figure 6. Discrimination net after the hypothesis that white square predicts Category A has been memorized.

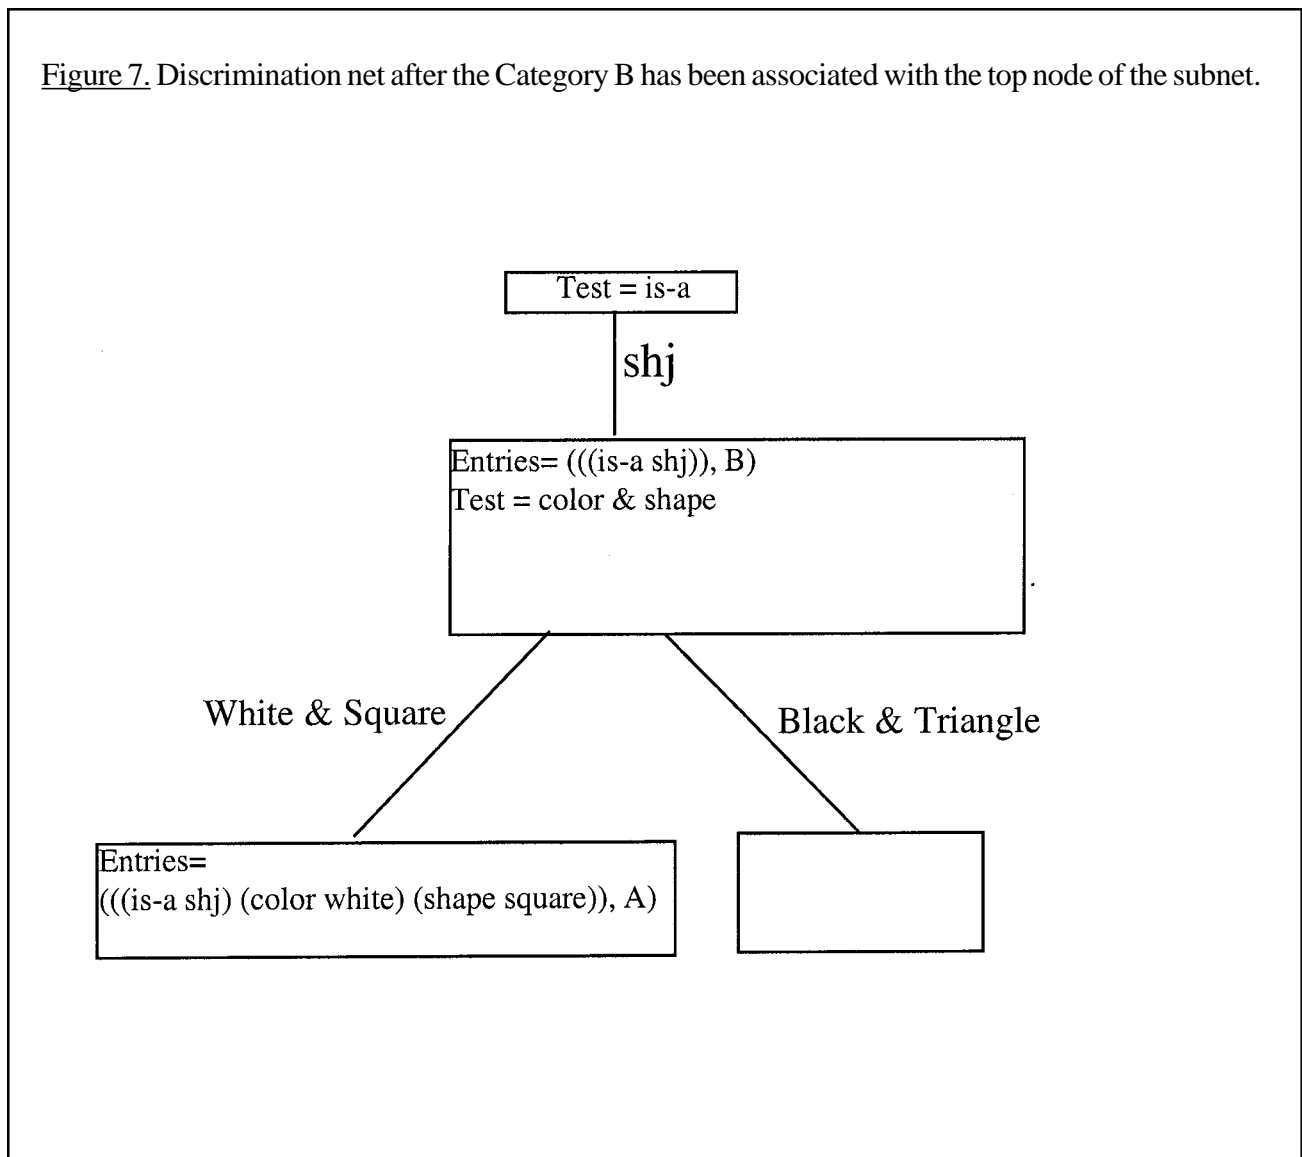


the discrimination net will correctly classify white squares using the bottom left node and all members of Category B using the top node of the “shj” subnet. The only items that it will not yet classify correctly are black triangles.

EPAM then completes the discrimination net when it tries to create an hypothesis by comparing the current stimulus, a large black triangle with the previous stimulus, a small white square, which was also a member of Category A. Since step 5 cannot find the features for an hypothesis, it instead causes the black triangle node to be associated with Category A. The discrimination net is now complete as shown in Figure 8.

Nosofsky, Gluck et al. (1994) also reported the time course of learning for human subjects doing the

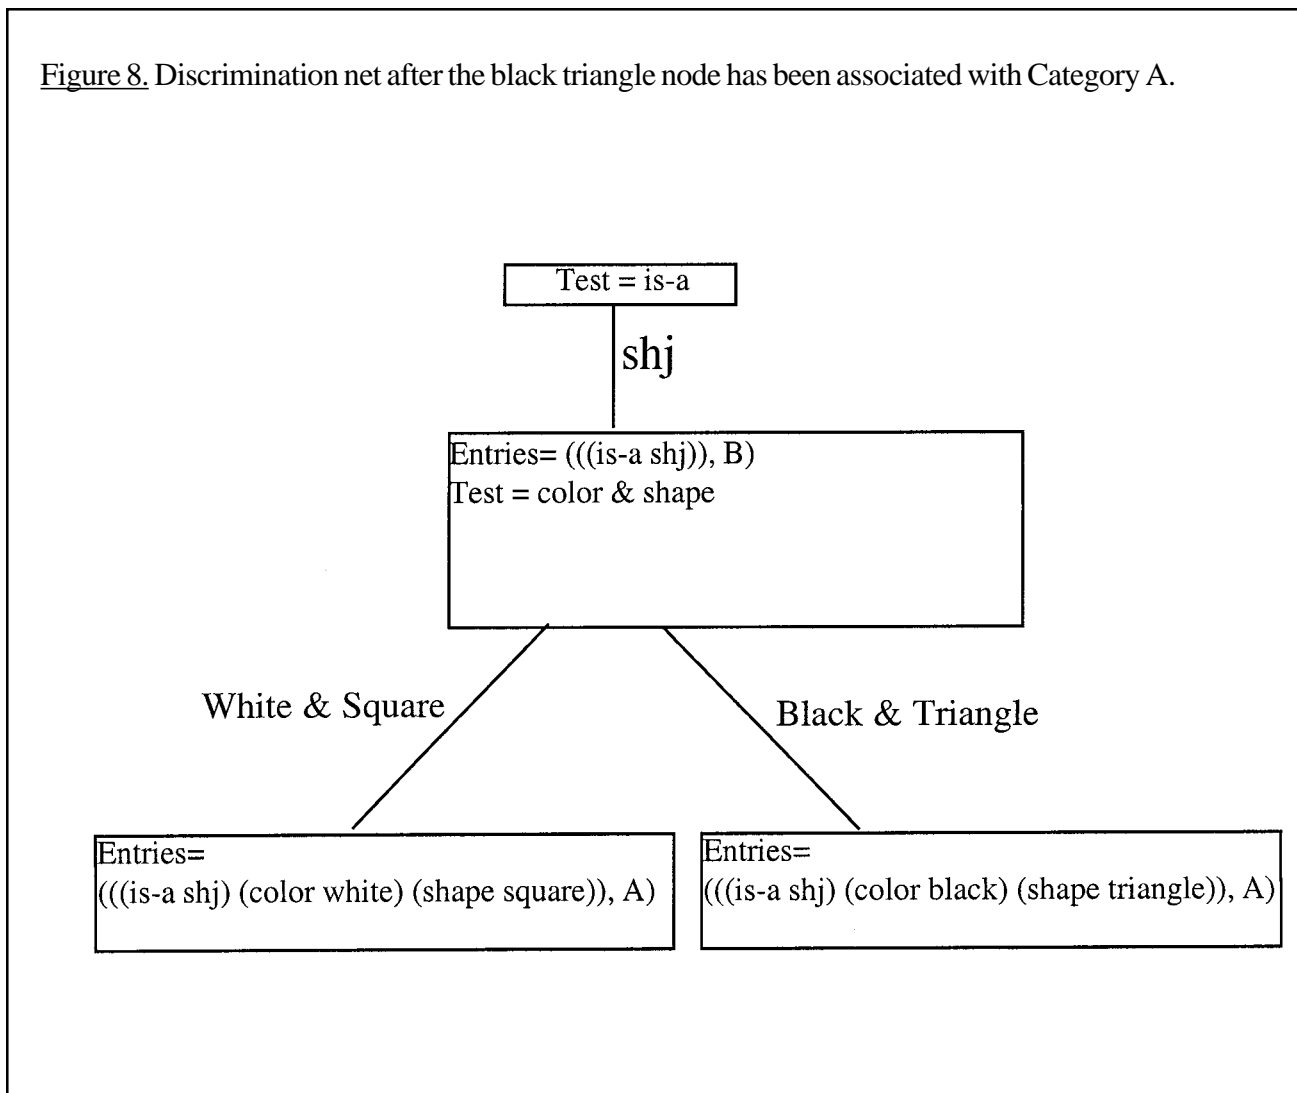
Figure 7. Discrimination net after the Category B has been associated with the top node of the subnet.



Type II problem and Nosofsky, Palmeri, and McKinley (1994) reported the same data for RULEX. A comparison between the data for human subjects, RULEX, and EPAM is shown in Table 2. All three learn the concept quickly, with the number of errors rapidly declining over the course of the first four blocks. By the fifth block, people are only making errors on 3% of the stimuli, RULEX is only making errors on about 5% of the stimuli, and EPAM is only making errors on 5% of the stimuli.

RULEX comes closer than EPAM to the human data. In order to get this good result one of RULEX's key free parameters ("branch") had to be changed from its default value. That parameter determines the ordering of the components of RULEX's rule-finding strategy. At its default setting, RULEX always looks for imperfect single dimensional rules before it searches for conjunctive two dimensional rules. However in its

Figure 8. Discrimination net after the black triangle node has been associated with Category A.



simulation of these six tasks, 90% of the time RULEX searches for conjunctive two dimensional rules before it searches for imperfect single dimensional rules. EPAM has no such parameter. Its strategy looks for perfect conjunctive rules at the same time that it looks for imperfect single dimensional rules.

It is not surprising that both EPAM and RULEX solve this problem fairly quickly. Although the strategies that they use differ in detail, both use strategies that search for either predictive single features or predictive conjunctions of two features.

Task III, IV, and V

Shepard, Hovland, and Jenkin's (1961) third, fourth and fifth tasks were all more difficult for their human subjects than their first two tasks, but less difficult than the sixth task. The overall results for EPAM in all six tasks are shown in Table 4. Nosofsky, Palmeri, and McKinley (1994) further analyze human subject

Table 2

Average Error Proportions for Type II problems, in Blocks 1-6 of Learning.

Block	People	RULEX	EPAM
1	0.38	0.42	0.45
2	0.16	0.12	0.31
3	0.08	0.08	0.21
4	0.06	0.07	0.11
5	0.03	0.05	0.05
6	0.03	0.04	0.02

Note. Human data is from Nosofsky, Gluck et. al. (1994). Rulex data is estimated from the graph shown as Figure 5 of Nosofsky, Palmeri, and McKinley (1994).

Table 3

Proportion of Errors for the Type V problem in Blocks 1-16 of Learning

Stimulus	People	RULEX	EPAM
Central	0.08	0.08	0.10
Peripheral	0.11	0.12	0.17
Exception	0.17	0.19	0.31

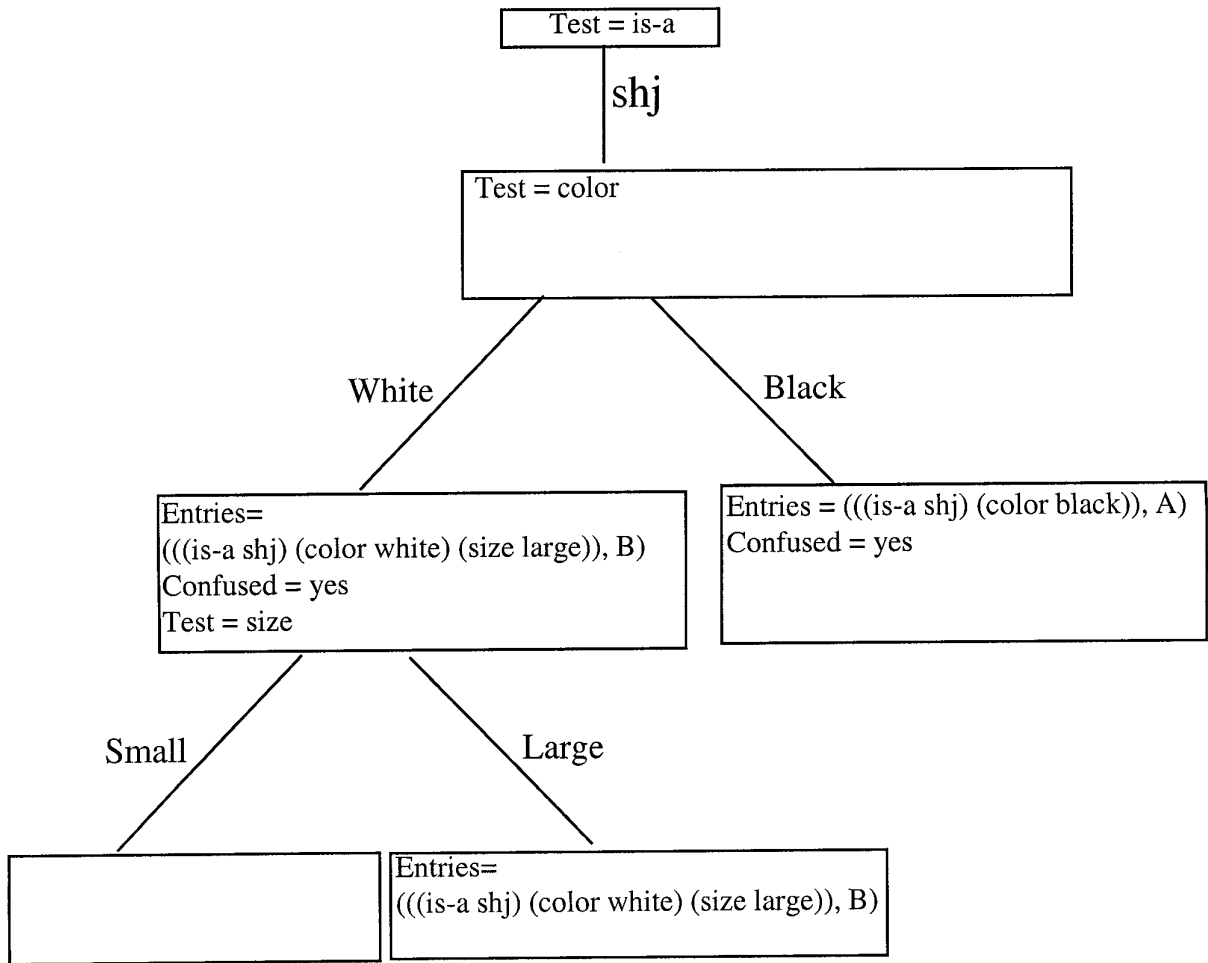
Note. Human data and RULEX data are estimated from Nosofsky, Gluck et al. (1994) Figure 9.

performance in these tasks by classifying stimuli as either central or peripheral, or as central, peripheral, and exception stimuli. For example in Task 5 where the members of Category A could be black & square or square & large or white & triangle & small, the large & black & square stimulus was considered to be central, the white & triangle & small stimulus was considered to be an exception, and the other two members of Category A were considered to be peripheral. EPAM, like RULEX, predicts the overall result shown with people that the central stimuli are learned most quickly, peripheral less quickly, and exception least quickly as shown in Table 3.

Here are some highlights from a simulation run with Task V which show why this is so. In this run, the central member of Category A was a large black triangle, the peripheral members were a small black triangle and a large black square, and the exception member was a small white square. The central member of Category B was a large white triangle, the peripheral members were a small white triangle and a large white square, and the exception member was a small black square.

The top test is a test for color as in the net shown in Figure 5. This test is chosen using the hypothesis testing procedures because three of the four members of Cat-

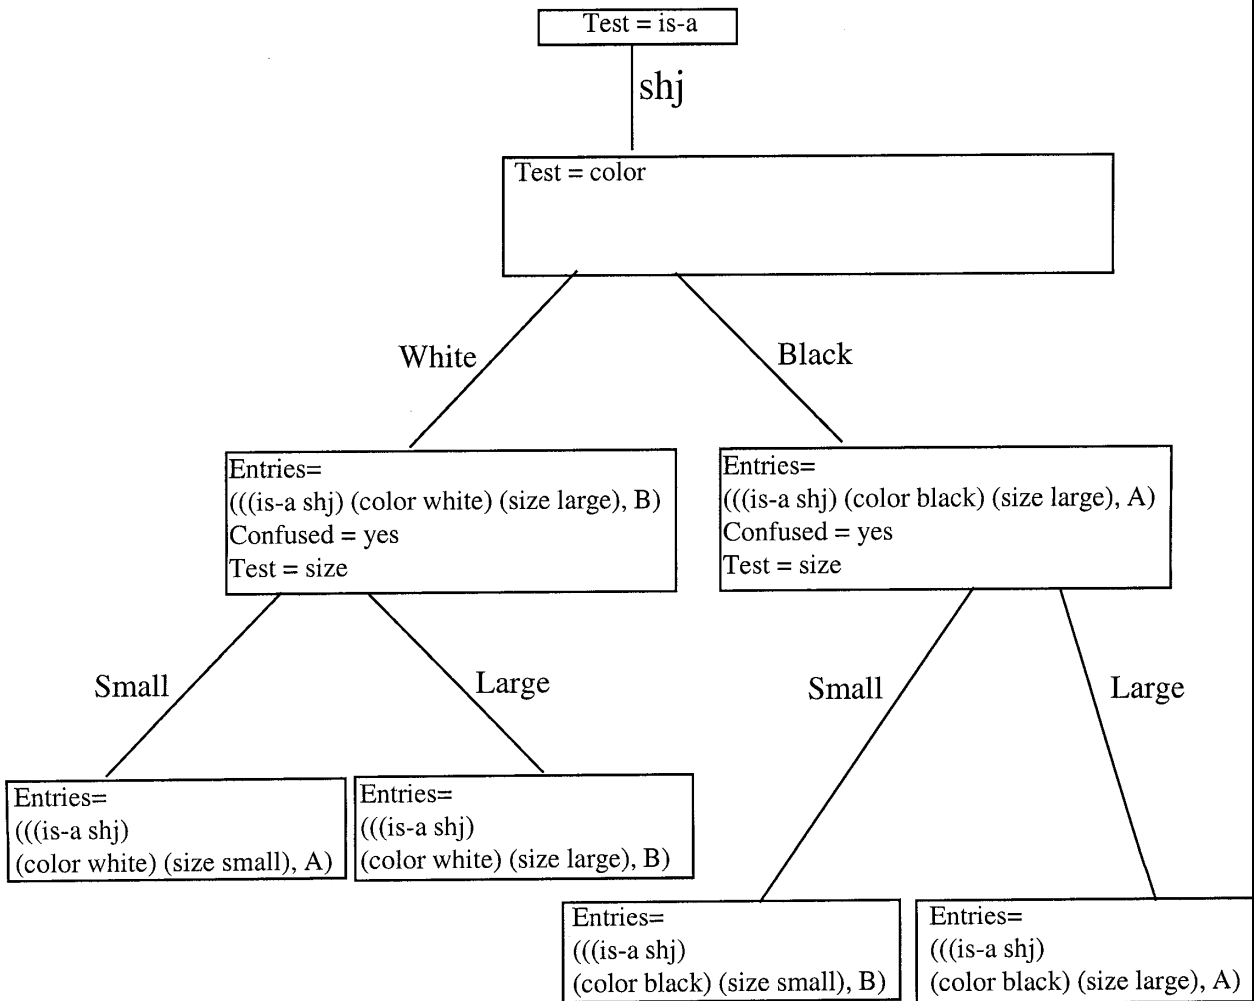
Figure 9. Discrimination net after the hypothesis that large and white predicts Category B has been memorized.



Category A are black and three of the four members of Category B are white.

After the hypothesis that large and white predicts Category B is added, the net shown in Figure 9 has been created.⁸ At this point EPAM will always be correctly categorizing the central member of Category B (the large white triangle) as well as one of the peripheral members of Category B (the large white square).


Figure 10. Discrimination net after the hypothesis that large & black predicts Category A has been memorized and after blank nodes have been associated with a category.



The next hypothesis to be added to the discrimination net is that large & black predicts Category A, which, after the blank nodes have been associated with a category, produces the discrimination net shown in Figure 10.

From here on the system will always be classifying all of the large black stimuli and all of the small white

Figure 11. Discrimination net at the end of a Task V experimental run.



stimuli. These consist of the central members of both categories as well as half of the peripheral members. The task remaining for EPAM is to discriminate the small black triangle (a peripheral member of Category A) from the small black square (the exception member of Category B) and the small white triangle (a peripheral member of Category B) from the small white square (the exception member of Category A). At the end of the experiment, the discrimination net is the one shown in Figure 11.

As we have illustrated by this example, EPAM predicts the difference in error rates for central, peripheral, and exception stimuli (shown in Table 3) because: (1) the central stimuli are sorted out by tests that are added to the discrimination net early in a simulation run; (2) some, but not all, of the peripheral stimuli are also sorted out by those early tests; and (3) the last tests added to the discrimination net discriminate an exception stimulus from a peripheral stimulus of the other category.

Task VI

The sixth task is clearly the hardest for both human subjects and EPAM. RULEX requires about the same amount of time as people to solve this problem. However, EPAM makes many more errors. In this task precisely two members of each category share each feature and no two members of any category share two features in common. As a result, EPAM typically searches fruitlessly for predictive features or conjunctions of two predictive features before it is finally able to accept a top test in the subnet. Almost as soon as an hypothesis is formed, it is rejected.

It is possible that the greater success of people when compared to EPAM on this task indicates that people tend to switch to a different strategy when they get frustrated. Perhaps they switch to a rote learning strategy such as the one programmed into EPAM for use when simulating the paired associate paradigm (Simon and Feigenbaum, 1964; Feigenbaum and Simon, 1984). The structure of EPAM's discrimination nets was worked out in order to fit the rich data of that paradigm.

Perhaps an additional parameter could be programmed into EPAM which would cause EPAM to switch to rote learning mode if a certain amount of time had gone by without progress toward a problem solution. This would actually cause EPAM to more closely resemble RULEX which first learns a rule (equivalent to the top node of an EPAM net) through a hypothesis-testing procedure but then switches to a rote learning mode in order to memorize the exceptions to the rule. The strategies of EPAM and RULEX could

Table 4

Average Error Proportions for Each of Problem Types 1-VI in Blocks 1-25 of Learning

Block	People						EPAM VI					
	I	II	III	IV	V	VI	I	II	III	IV	V	VI
1	0.21	0.38	0.46	0.42	0.47	0.50	0.24	0.45	0.44	0.44	0.48	0.58
2	0.03	0.16	0.29	0.30	0.33	0.34	0.06	0.31	0.36	0.37	0.40	0.56
3	0.00	0.08	0.22	0.22	0.23	0.28	0.00	0.21	0.30	0.31	0.36	0.55
4	0.00	0.06	0.15	0.17	0.14	0.25	0.00	0.11	0.26	0.27	0.32	0.54
5	0.00	0.03	0.08	0.15	0.11	0.22	0.00	0.05	0.24	0.23	0.28	0.53
6	0.00	0.03	0.08	0.11	0.08	0.19	0.00	0.02	0.21	0.20	0.25	0.52
7	0.00	0.03	0.06	0.09	0.07	0.19	0.00	0.01	0.17	0.16	0.21	0.50
8	0.00	0.02	0.03	0.06	0.08	0.18	0.00	0.00	0.14	0.13	0.18	0.48
9	0.00	0.02	0.02	0.03	0.05	0.17	0.00	0.00	0.11	0.10	0.15	0.47
10	0.00	0.01	0.02	0.03	0.05	0.13	0.00	0.00	0.08	0.08	0.12	0.44
11	0.00	0.00	0.02	0.02	0.05	0.14	0.00	0.00	0.05	0.05	0.08	0.42
12	0.00	0.00	0.01	0.03	0.04	0.12	0.00	0.00	0.03	0.03	0.06	0.41
13	0.00	0.01	0.01	0.01	0.03	0.10	0.00	0.00	0.02	0.02	0.04	0.39
14	0.00	0.00	0.01	0.00	0.03	0.10	0.00	0.00	0.01	0.01	0.03	0.36
15	0.00	0.00	0.01	0.00	0.02	0.11	0.00	0.00	0.00	0.00	0.02	0.34
16	0.00	0.00	0.01	0.00	0.01	0.11	0.00	0.00	0.00	0.00	0.01	0.32
17	0.00	0.00	0.01	0.00	0.01	0.08	0.00	0.00	0.00	0.00	0.01	0.29
18	0.00	0.00	0.01	0.00	0.01	0.08	0.00	0.00	0.00	0.00	0.00	0.27
19	0.00	0.00	0.01	0.00	0.01	0.08	0.00	0.00	0.00	0.00	0.00	0.25
20	0.00	0.00	0.00	0.00	0.01	0.06	0.00	0.00	0.00	0.00	0.00	0.24
21	0.00	0.00	0.01	0.00	0.01	0.06	0.00	0.00	0.00	0.00	0.00	0.21
22	0.00	0.00	0.00	0.00	0.01	0.04	0.00	0.00	0.00	0.00	0.00	0.20
23	0.00	0.00	0.00	0.00	0.01	0.04	0.00	0.00	0.00	0.00	0.00	0.18
24	0.00	0.00	0.01	0.00	0.01	0.03	0.00	0.00	0.00	0.00	0.00	0.17
25	0.00	0.00	0.00	0.00	0.01	0.04	0.00	0.00	0.00	0.00	0.00	0.16
Ave	0.01	0.03	0.06	0.07	0.08	0.14	0.01	0.05	0.10	0.10	0.12	0.38

Note. Human data is from Nosofsky, Gluck et. al. (1994)

merge to a point where the only difference between the two would be EPAM's use of a discrimination net.

Overall Results for the Six Tasks

Shepard, Hovland & Jenkin's basic finding, also found in a replication conducted by Nosofsky, Gluck et al. (1994), was that human subjects found the Type I task to be the easiest, followed by the Type II task. Tasks III, IV, and V were similarly difficult. And Task VI was the most difficult. EPAM finds the same pattern. The bottom row of Table 4 shows the average error proportion or errors per task for people and for EPAM over the 25 blocks of the experiment. The average proportion of errors per task for EPAM (.01, .05, .10, .10, .12, and .38) shares the same ordering as the average proportion of errors per task for Nosofsky, Gluck et al.'s human subjects (.01, .03, .06, .07, .08, and .14) though the error rates for EPAM tend to be higher than those of people..

EPAM especially does worse with the Type VI task than people perhaps because this particular task is not amenable to hypothesis formation. If EPAM could be programmed with a parameter that would switch to a rote learning strategy when hypothesis formation has been repeatedly failing, EPAM might fit the human data more closely.

Martin and Caramazza Tasks

In studies of defining-features and family-resemblance categories, Martin and Caramazza (1980) not only collected statistics about their subjects' performance, but also recorded subjects' comments and carefully analyzed how their subjects were doing the tasks. Both experiments classified stimuli into two categories called "Harry" and "Charlie." Figure 12 illustrates the features used in these experiments. We will discuss each experiment in turn.


EPAM's simulations of both of Martin and Caramazza's experiments were run with 1000 simulated subjects and the following free parameters: (1) speed-of-learning-parameter = 100, (2) weight for hair = 10, (3) weight for face-shape = 22, (4) weight for eyes = 4, (5) weight for eye-brows = 1, (6) weight for nose = 1, (7) weight for mouth = 2, and (8) weight for ears = 6.

Defining Features Experiment

In their defining features experiment (Experiment 2), Martin and Caramazza used the categories shown in Table 5. Here is how Martin and Caramazza describe this condition:

For the defining features condition, the categories were structured so that a conjunctive rule

Figure 12. Faces showing examples of all the features used in Martin and Caramazza's experiments. [Note. This is Figure 1 from Martin & Caramazza, 1980.]



would define category membership: The members of Category 1, referred to as Harrys, had round noses and hair whereas the members of Category 2, referred to as Charlies, had frowns and eyes looking to the left. The remaining features were more or less typical of one category or the other. For example, of the 12 Harrys, 9 had ears and 7 had thin faces, whereas of the 12 Charlies, 9 had no ears and 7 had fat faces. In order that a conjunctive rule would be necessary for determining category membership, some members of the contrasting category had one of the defining features of the other category (but never both). For example, 4 of the Charlies had round noses and 4 had hair, but none had both a round nose and hair. (p. 327)

Martin and Caramazza collected information about response latencies, trial of last error (TLE), and typicality ratings and reported correlations between their results and the predictions of various models, including Medin and Schaeffer's (1978) context model, that had been developed in order to explain family-resemblance data. They reported that these models did not do very well in this defined-features task: "Although

Table 5

Structure of Defining Features Categories

Category/ picture number	Hair ^a	Face- shape ^b	Face- eyes ^c	Eye brows ^d	Nose ^e	Mouth ^f	Ears ^g
Harry							
1	1	2	1	1	1	1	1
2	1	1	2	1	1	1	1
3	1	1	1	2	1	1	1
4	1	2	2	2	1	1	1
5	1	1	2	2	1	1	1
6	1	1	2	2	1	1	2
7	1	2	1	1	1	2	1
8	1	1	1	2	1	2	1
9	1	2	1	2	1	2	1
10	1	1	1	1	1	2	1
11	1	2	1	2	1	2	2
12	1	1	1	2	1	2	2
No. of 1s	12	7	9	4	12	6	9
Charlie							
1	2	2	2	2	1	2	2
2	2	1	2	2	1	2	2
3	2	2	2	1	1	2	1
4	2	1	2	1	1	2	2
5	2	2	2	1	2	2	2
6	1	2	2	2	2	2	2
7	2	1	2	1	2	2	2
8	1	2	2	1	2	2	2
9	2	2	2	2	2	2	2
10	1	1	2	1	2	2	2
11	2	2	2	1	2	2	1
12	1	1	2	1	2	2	1
No. of 1s	4	5	0	8	4	0	3

Note. ^a1 = hair, 2 = no hair. ^b1 = thin face, 2 = fat face. ^c1 = eyes left, 2 = eyes up. ^d1 = curved eye-brows, 2 = straight eyebrows. ^e1 = round nose, 2 = triangular nose. ^f1 = smile, 2 = frown. ^g1 = ears, 2 = no ears. This table is reproduced from Table 1 of Martin and Caramazza (1980).

some of the models had high correlations with reaction times for one category, none did a good job of predicting reaction times for both categories” (p. 334).

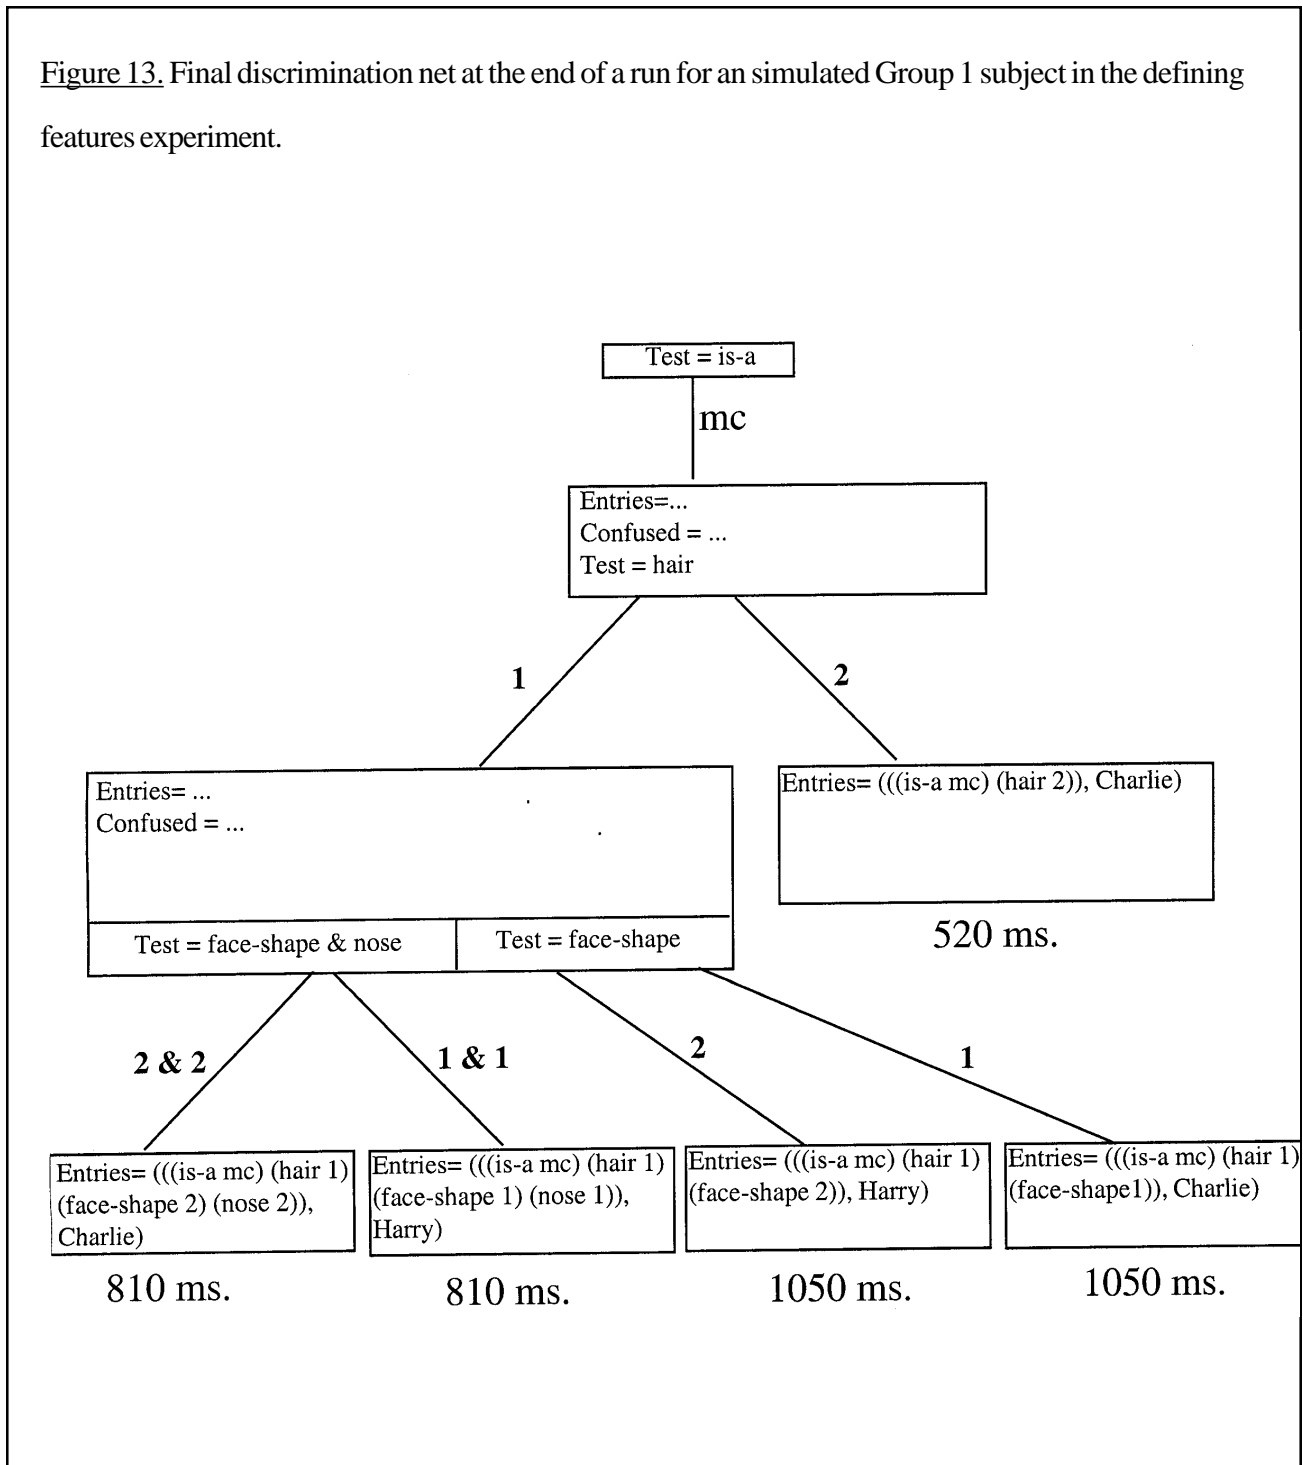
When they evaluated individual differences quantitatively using an inverse principal-components analysis on the response latencies they found that latency data divided their subjects into two groups. They also reported qualitatively different subject comments between the members of these two groups. All 12 of Martin and Caramazza’s subjects reached their criterion of two consecutive perfect trials within the hour-long training period; 90% of EPAM’s simulated subjects reached that criterion.

Group 1. According to subject comments, the 6 of the 12 subjects who were in Group 1 tested for either nose or hair while 40% of the simulated EPAM subjects who reached criterion tested first for hair and an additional 3% tested first for nose. Martin and Caramazza reported that their subjects in this group had longer latency response times for those Charlies that had either round noses or hair compared to other Charlies. The same was true for EPAM: Charlies with either round noses or hair had average latencies of 0.96 seconds compared to other Charlies at 0.75 seconds.

The final discrimination net at the end of an EPAM run for a simulated subject in this group, shown in Figure 13, shows why a test for hair would lead to longer reaction times for Charlies with hair. The response latencies for this simulated subject are also shown in the figure. Each latency is essentially a function of the number of tests traversed in the discrimination net when classifying the stimulus with each test requiring approximately one eye fixation (250 ms.) of time⁹. In this run, the Charlies with hair (i.e. hair = 1) took either 810 ms. or 1050 ms. to classify which was longer than the 520 ms. that it took to classify the Charlies without hair (i.e. hair = 2). That was because the top test was a test for hair and only the Charlies without hair (i.e. hair = 2) were discriminated by just a single test.¹⁰

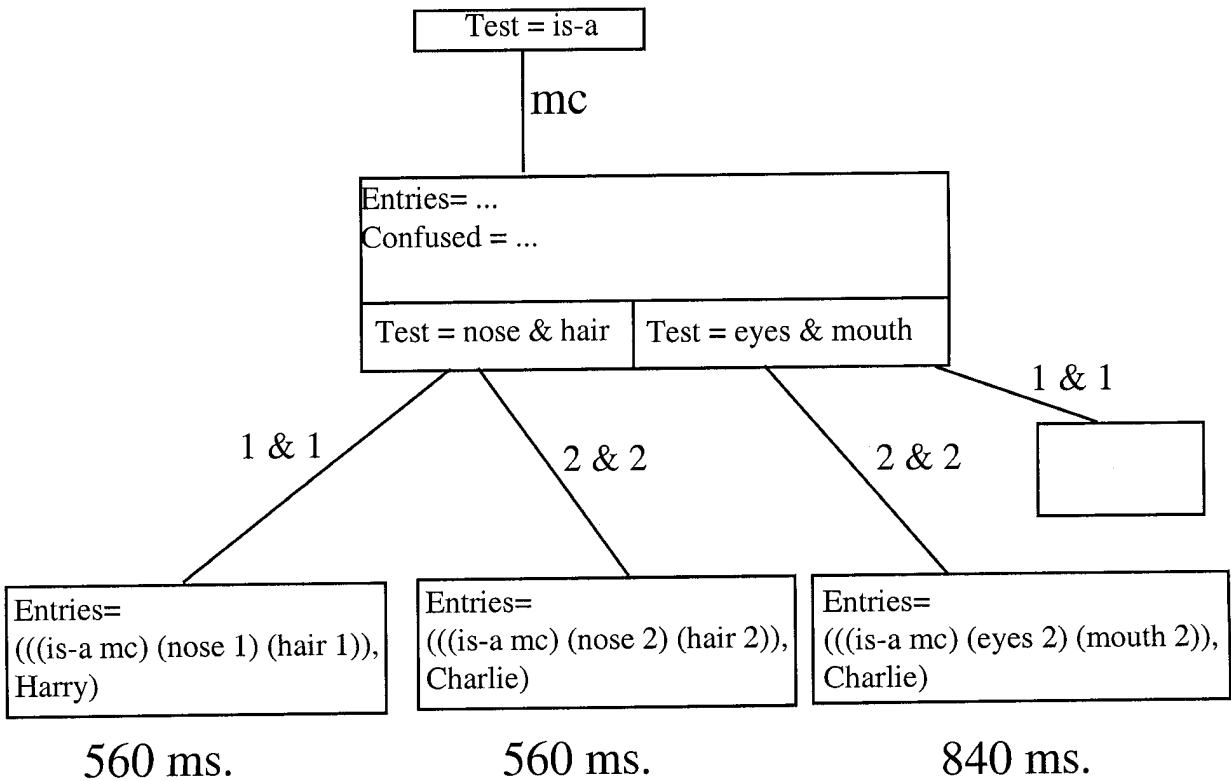
Group 2. According to subject comments the 5 of the 12 subjects who were in Group 2 learned a conjunction of nose and hair, the defining characteristics of the Harry category and then classified the other stimuli as Charlies; 16% of the simulated subjects who reached criterion in the EPAM runs were in this group. Figure 14 shows a typical discrimination net with latency times that is grown during an experimental run for a simulated subject in this group. Martin and Caramazza report that the average latency of response time for Harrys was 0.68 seconds and the average latency of response time for Charlies was 0.87 seconds. EPAM also shows a shorter latency time for Harrys (0.56 seconds) compared to Charlies (0.94 seconds). This

Figure 13. Final discrimination net at the end of a run for an simulated Group 1 subject in the defining features experiment.



occurs because Harrys are discriminated by a single test, a test for nose & hair, while some Charlies are first sorted through a test for nose & hair and then, because they match neither of the values that are in the discrimination net, they are then sorted through a second test, in the case of the run shown in Figure 14, the second test is a test for the conjunction eyes & mouth.

Figure 14. Final discrimination net at the end of a run for an simulated Group 2 subject in the defining features experiment.



Unclassified Subjects. While only 1 of 12 of Martin and Caramazza’s subjects was unclassifiable into Group 1 or Group 2 by their analysis, 41% of EPAM’s simulated subjects who reached criterion did not fit Martin and Caramazza’s description of the discrimination nets used by Groups 1 and 2. Of the remaining simulated subjects, 11% had a test for eyes at the top of their discrimination net (the Charlie category counterpoint to hair), 10% had a conjunctive test for eyes & mouth (the Charlie category counterpoint to hair & nose) and the other 20% had various other tests at the top of their discrimination nets.

It appears that people were much more likely than EPAM to put “hair” into the test at the top node of their discrimination nets. Perhaps the similarity between the word “Harry” and the word “hair” caused Martin

and Caramazza's subjects to form hypotheses that included hair at the beginning of the experiment. We have tried increasing the feature weight of hair but that throws off results in the family-resemblance condition. It may be that hair is only especially salient at the very beginning of the experiment.

Family Resemblance Experiment

Here is how Martin and Caramazza described the family-resemblance categories, shown in Table 6 that they used in their third experiment:

The two categories were symmetric with respect to each feature, that is, if one category had n faces with a certain feature, the other category had $12-n$ faces with this feature. For example, 9 of the 12 Harrys had ears, so 3 of the Charlies had ears. The categories were also structured so that each face would have a greater family resemblance to its own category than to the other category, with family resemblance being determined in terms of feature overlap... (p. 339)

In a trial run with three subjects, Martin and Caramazza found that they had to change their category structures in this condition. At first they had assigned eyes and ears to be the most predictive, but their three trial run subjects were still making 6 to 8 errors on the 24 faces after a full hour of testing. In the revised category structure they made most characteristic what they described as the most salient features (mouth and face shape).

Martin and Caramazza reported that 8 of their 17 human subjects failed to meet their criteria of two consecutive perfect trials through the list within one hour. Similarly 41% of EPAM's simulated subjects failed to reach the criterion of two consecutive perfect trials during the maximum of 30 trials that we presented to each simulated subject.

Martin and Caramazza divided the remaining 9 subjects into two groups through an inverse principal components analysis and were then able to deduce discrimination nets of each group showing subject's latency times as a function of the tests traversed in the discrimination nets. Figure 15 shows the deduced discrimination net and latency times for their 4 subjects in Group 1 and Figure 16 shows the deduced discrimination net and latency times for their 5 subjects in Group 2.

When run as simulated subjects in this experiment, 21% of EPAM's simulated subjects who reached criteria were in Group 1, 64% were in Group 2, and 15% were unclassifiable as Group 1 or Group 2. If the top


Table 6

Structure of Revised Family Resemblance Categories

Category/ picture number	Hair ^a	Face- shape ^b	Face- eyes ^c	Eye brows ^d	Nose ^e	Mouth ^f	Ears ^g
Harry							
1	1	2	1	1	1	1	1
2	1	2	2	2	2	1	1
3	1	2	1	1	1	2	1
4	1	1	1	2	2	1	1
5	2	2	1	2	2	1	1
6	2	1	1	1	1	1	1
7	2	2	2	2	1	2	2
8	2	2	2	1	1	2	1
9	2	2	2	1	2	1	2
10	1	2	2	2	2	1	2
11	2	2	1	2	2	1	2
12	2	2	2	1	2	2	1
No. of 1s	5	2	6	6	5	8	8
Charlie							
1	2	1	2	1	1	1	1
2	1	1	1	2	2	2	1
3	2	1	2	2	1	1	1
4	1	1	2	2	1	1	2
5	1	2	1	2	1	2	2
6	1	1	1	2	1	2	2
7	2	1	2	2	1	2	2
8	2	1	1	1	1	2	2
9	2	1	2	1	2	1	2
10	1	2	1	1	2	2	2
11	1	1	1	1	2	2	2
12	1	1	2	1	2	2	1
No. of 1s	7	10	6	6	7	4	4


Note. ^a1 = hair, 2 = no hair. ^b1 = thin face, 2 = fat face. ^c1 = eyes left, 2 = eyes up. ^d1 = curved eyebrows, 2 = straight eyebrows. ^e1 = round nose, 2 = triangular nose. ^f1 = smile, 2 = frown. ^g1 = ears, 2 = no ears. This table is reproduced from Table 4 of Martin and Caramazza (1980).

Figure 15. Discrimination net and latency times deduced by Martin and Caramazza's (1980) for their Group 1 subjects in the family resemblance experiment. [Note. This is Figure 7 from Martin & Caramazza, 1980.]



test was “face-shape” and the top-left test was “hair” (as shown in Figure 15) then we put the subject into Group 1. We also put those subjects whose top test was “hair” or those whose top test was a conjunction of “face-shape” and “hair” into Group 1. We put any other subject with a top test for “face-shape” into Group 2.

Figure 16. Discrimination net and latency times deduced by Martin and Caramazza's (1980) for their Group 2 subjects in the family resemblance experiment. [Note. This is Figure 8 from Martin & Caramazza, 1980.]



We only considered a subject to be unclassifiable if the top test was neither face-shape, nor hair, nor a conjunction of the two.

EPAM does a good job of fitting the latency time data of Figures 15 and 16 as shown in Tables 7 and 8 which compare the latency times of people with those of EPAM. EPAM's ordering of the latency times in each chart is the same as the ordering for Martin and Caramazza's subjects. EPAM assumes that it takes approximately 250 ms. to sort through each test in the discrimination net that requires an eye fixation. Its good fit with the human data shows that the 85% of EPAM's simulated subjects who either put "face shape" or

Table 7

Response latency times corresponding to the discrimination net shown in Figure 15

Group of Faces	People	EPAM
Harry 5,7,8,9,11,12	789	770
Harry 1,2,10	996	1167
Harry 3	1257	1297
Charlie 5,10	1322	1507
Charlie 4,6,7,8,9,11	731	909
Charlie 1,3,12	939	1089
Charlie 2	1226	1203
Harry 4,6	1276	1304

Note. Human data is from Martin and Caramazza (1980). All latency times are expressed in milliseconds.

Table 8

Response latency times corresponding to the discrimination net shown in Figure 16

Group of Faces	People	EPAM
Harry 1,2,3,5,8,9,10,11,12	819	960
Harry 7	993	1196
Charlie 5,10	1271	1519
Charlie 2,4,6,7,8,9,11,12	867	960
Charlie 1,3	1006	1075
Harry 4,6	1117	1252

Note. Human data is from Martin and Caramazza (1980).

All latency times are expressed in milliseconds.

“hair” or a combination of the two into the top node of the discrimination net produced discrimination nets that were very similar to those produced by Martin and Caramazza’s subjects.

One reason for the good fit is due to the fact that EPAM so often puts face shape at the top of its nets. This is partly due to the fact that face shape, as shown in the third column of Table 6, is one of the most predictive features in that 10 of the 12 members of the Harry category have thin faces while 10 of the 12 Charlies have fat faces. The other reason why face shape was chosen so often by simulated subjects was that its weighting was set higher by the experimenter than the other features. If not for this weighting, face shape would not as often appear in the hypothesis and so would not as often be chosen as the top test in the net. Yet even this weighting is not entirely a free parameter: Martin and Caramazza reported (p. 339) that mouth and face shape were “the most salient features” and that they had purposely assigned salient features to be predictive so that subjects would notice them.

But there is much more to the net than face shape. EPAM predicts the same nets as Martin and Caramazza's subjects right down the line. The reason for this good fit is EPAM's strategy, designed with means-ends analysis in mind, which searches for tests that allow the number of unclassifiable items to be reduced. Those are exactly the sort of tests that Martin and Caramazza's subjects adopted.

For example, on the left side of Figure 15 these Group 1 subjects first memorized a test for "hair" which allowed them to correctly classify Harrys 5, 7, 8, 9, 11, and 12. Then they memorized a test for "mouth" which allowed them to classify Harrys 1, 2, and 10. The same thing happened on the right side. The test for "ears" allowed them to classify Charlies 4, 6, 7, 8, 9, and 11 and the test for eyes allowed them to classify Charlies 1, 3, and 12. Every test that they memorized allowed them to reduce the number of unclassifiable items. In hindsight, EPAM's strategy could have been deduced simply from this figure and Figure 16.

RULEX would not put together these discrimination nets. Though it would probably adopt the test for face-shape as its rule, it would not then proceed, like EPAM, to search for additional tests that would reduce the number of unclassifiable items. Thus it would not be able to capture the response latency data.

Although EPAM's fit with the human data in the family resemblance experiment is excellent, its fit with the data of the defined features experiment could be improved if EPAM were to notice the similarity between the attribute "hair" and the Category "Harry." As currently programmed, however, EPAM does not notice such acoustic/semantic similarities.

The 5-4 Task Randomized

Medin and Schaeffer's (1978) four-attribute binary-value family-resemblance category structure, sometimes called the 5-4 structure, is one of the most extensively studied classification structures ever (Smith and Minda, 2000). Here is Nosofsky, Palmeri, and McKinley's (1994) description of that structure (see the "values" column of Table 9):

The stimuli vary along four binary-valued attributes. There are five training exemplars in Category A and four in Category B; the remaining seven stimuli are transfer items that are not assigned by the experimenter to either category. In general, members of Category A tend to have a logical value of 1 on each of their attributes, whereas members of Category B tend to have a logical value of 2. Importantly, however, no single-attribute rule is available for perfectly partitioning the exemplars into categories, and no conjunctive rule is available either. . . . (pp.

Table 9

Categoryization Response Probabilities of reporting that item is a member of
Category A during the Transfer Condition

Face Values	People	RULEX	Context	EPAM
Category A				
4 (1112)	0.77	0.79	0.79	0.83
7 (1212)	0.78	0.83	0.79	0.84
15 (1211)	0.83	0.88	0.77	0.90
13 (1121)	0.64	0.65	0.65	0.63
5 (2111)	0.61	0.64	0.63	0.63
Category B				
12 (1122)	0.39	0.45	0.40	0.47
2 (2112)	0.41	0.44	0.40	0.47
14 (2221)	0.21	0.23	0.21	0.28
10 (2222)	0.15	0.16	0.19	0.20
Transfer Item				
1 (1221)	0.56	0.62	0.58	0.57
3 (1222)	0.41	0.47	0.47	0.46
6 (1111)	0.82	0.85	0.79	0.91
8 (2212)	0.40	0.45	0.45	0.47
9 (2121)	0.32	0.34	0.33	0.39
11 (2211)	0.53	0.61	0.56	0.58
16 (2122)	0.20	0.22	0.22	0.26

Note. The Human data and the RULEX data are from Nosofsky, Palmeri and McKinley (1994).

55-56)

In their first experiment Nosofsky, Palmeri, and McKinley (1994) replicated Medin and Schaeffer's (1978) task except that they randomly assigned physical categories to logical categories allowing models like RULEX and EPAM to equalize the attribute weights in their simulations. EPAM's simulation of this experiment was run with 10,000 simulated subjects and the speed-of-learning-parameter set at 40.

RULEX, EPAM, and the context model (Medin & Schaeffer, 1978) capture the main results of this experiment quite well. Table 9 shows probability of classifying a stimulus as Category A during the transfer phase of the experiment. Although both RULEX and the context model capture the results a bit more closely than EPAM, there is little difference between the human results, the RULEX results, and the context model's results.

Nosofsky, Palmeri, and McKinley also compared how well RULEX and the context model could account for the observed distribution of generalizations observed at the subject level. Here is how Nosofsky, Palmeri, and McKinley (1994) describe this generalization data:

Because there were 7 transfer stimuli and 2 categories, there are $2^7 = 128$ possible patterns of generalization, although only 36 of these patterns were exhibited by at least 2 subjects. For example, the pattern AAABBBB, which was exhibited by 32 subjects, corresponds to those subjects who classified Transfer Stimuli 1 to 3 in Category A and Transfer Stimuli 4 to 7 in Category B. (Because there were three transfer blocks, a subject is said to classify a stimulus into Category A during the transfer phase if he or she classifies it into Category A in at least two of the three blocks.) Two of the most common generalizations were AAABBBB and BBAABAB. The former is consistent with the Attribute 1 rule: $1*** \rightarrow A, 2*** \rightarrow B$; the latter is consistent with the Attribute 3 rule $**1* \rightarrow A, **2* \rightarrow B$. A third common generalization was ABABBAB. (p.71)

As shown in Table 10, EPAM shows some of the same peek generalization patterns as people and RULEX. For example, the AAABBBB pattern occurs 14.1% of the time for people, 12.2% of the time for RULEX, 9.9% of the time for EPAM and just 1.5% of the time for the context model. Similarly the BBAABAB pattern occurs 13.2% of the time for people, 12.6% of the time for RULEX, 10.4% of the time for EPAM and just 6.4% of the time for the context model. The AAABBBB discrimination net usually has a test for the first

Table 10.
Generalization Pattern Probabilities during transfer stage

Pattern	People	Rulex	Context	EPAM
AAAAAAA	0.009	0.003	0.003	0.068
AAAABAB	0.013	0.012	0.033	0.023
AAABAAB	0.013	0.011	0.012	0.022
AAABABB	0.013	0.036	0.006	0.015
AAABBAB	0.026	0.034	0.031	0.020
AAABBBB	0.013	0.012	0.004	0.007
AAABBBB	0.141	0.122	0.015	0.099
AABBABA	0.009	0.004	0.000	0.002
AABBBBB	0.013	0.013	0.001	0.003
ABAAAAB	0.018	0.011	0.025	0.022
ABAABAB	0.022	0.034	0.062	0.021
ABABAAA	0.009	0.004	0.005	0.008
ABABAAB	0.026	0.015	0.024	0.037
ABABABA	0.009	0.005	0.003	0.004
ABABABB	0.009	0.014	0.012	0.008
ABABBAB	0.069	0.029	0.058	0.018
ABABBBB	0.026	0.035	0.029	0.033
ABBABAB	0.013	0.006	0.005	0.000
ABBBBAB	0.009	0.003	0.001	0.001
ABBBBBB	0.009	0.003	0.000	0.000
BAAABAB	0.013	0.012	0.034	0.010
BAABBAB	0.009	0.009	0.031	0.004
BAABBBB	0.013	0.012	0.016	0.016
BABBBBB	0.009	0.003	0.001	0.004
BBAAAAA	0.013	0.007	0.005	0.011
BBAAAAB	0.031	0.035	0.026	0.017
BBAABAB	0.132	0.126	0.064	0.104
BBAABBB	0.013	0.011	0.032	0.017
BBABAAA	0.018	0.004	0.005	0.004
BBABAAB	0.018	0.017	0.024	0.010
BBABABA	0.009	0.004	0.003	0.005
BBABABB	0.031	0.006	0.012	0.008
BBABBAB	0.035	0.035	0.059	0.036
BBABBBB	0.009	0.004	0.006	0.003
BBABBBB	0.044	0.009	0.030	0.045
BBBAABB	0.009	0.004	0.001	0.001

Note. The Human data, context model data, and RULEX data are estimated from Figures 9, 10, and 11 of Nosofsky, Palmeri and McKinley (1994).

dimension at the top and no other tests yet added to the discrimination net. The BBAABAB discrimination net usually has the third dimension at the top and no other tests yet added to the discrimination net.

Nosofsky, Palmeri, & McKinley note that the context model does a better job than RULEX with predicting the frequency of the ABABBAB generalization. The context model is also much better than EPAM with this particular generalization; 6.9% of the human subjects and 5.8% of the context subjects display this generalization compared to just 2.9% of the RULEX subjects and 1.8% of the EPAM subjects. In unreported data from an experiment conducted by Gluck et. al. (2001)¹¹ one of the human subjects (MS) found this very generalization in the very first trial (the first time the 9 faces were presented) of a classification learning experiment using this 5-4 task and as a result she was the only subject able to identify all 9 faces correctly during the second trial (the second time the 9 faces were presented). She reported using a “majority rule”

Table 11.

Hamming distance probabilities between first and second blocks during transfer stage

Hamming Distance	People	Rulex	Context	EPAM
0	0.26	0.26	0.03	0.36
1	0.25	0.31	0.14	0.25
2	0.23	0.18	0.28	0.21
3	0.16	0.11	0.29	0.12
4	0.09	0.08	0.18	0.05
5	0.01	0.05	0.06	0.01
6	0.01	0.02	0.01	0.00
7	0.00	0.00	0.00	0.00

Note. The Human data, RULEX data, and context model data are estimated from Figure 12 of Nosofsky, Palmeri & McKinley (1994)

strategy. If an item shared 2 of the three predictive features (Dimensions 1, 3, and 4) of Category A she classified it as Category A; if not then she classified it as Category B. It is apparent that strategies other than the ones used by RULEX and EPAM can be used by human subjects in a classification learning experiment.

EPAM especially overpredicts the AAAAAAA generalization, 6.8% of EPAM's simulated subjects used this generalization while only 0.9% of people used this generalization. This could either indicate a failing of the EPAM model or simply that people would be too embarrassed to respond with Category A to every stimulus for fear that the experimenter would think that they were not trying.

Nosofsky, Palmeri and McKinley also reported a measure of the constancy of the subjects' classifications. This measure, called "ham-

ming distances" counts how many of the seven stimuli an individual subject classifies differently during the second block of transfer trials than on the first. The results for people, RULEX, the context model, and EPAM are shown in Table 11.

The hamming distances for people, RULEX, and EPAM are skewed to the left while the hamming distances for the context model are distributed more normally with a mean between 2 and 3. Apparently EPAM, RULEX and people are more constant in their classifications during this transfer task than are the simulated individual subjects of the context model. RULEX changes its classifications as the result of its response error parameter which gives "the probability that a subject makes the opposite response indicated by the extracted rules" (p. 58). Normally this parameter is set at its default level of 0, but in this experiment it was

set at a level designed to give RULEX the same amount of variation in response as people.

EPAM only changes its classification of a stimulus from one transfer block to the next when it is guessing the classification. This occurs when the stimulus sorts to a discrimination net node that has not yet been associated with a category.

Although both EPAM and RULEX produce very good fits with the hamming distance data, RULEX relies upon a non-default setting of a special parameter to produce the proper amount of variability while EPAM produces this variability without any special parameters.

The 5-4 Task with Different Strategies

Medin and Smith (1981) conducted a classification learning experiment using the same 5-4 category structure as the Nosofsky, Palmeri, & McKinley (1994) experiment that we just discussed. While there were several minor differences (stimuli were Brunswik faces instead of rocket ships; assignment of physical to logical categories was not randomized), the main difference was in the experimental instructions. Each of three groups of subjects was given a different set of instructions: (1) in the standard condition subjects were given no specific instructions about how to do the categorization task, (2) in the rules-plus-exceptions condition the subjects were told to learn a Dimension 3 rule (nose-length), and memorize exceptions to that rule, and (3) in the prototype condition, subjects were told to learn the characteristics of Category A stimuli and those of Category B stimuli and then compare each stimulus with both prototypes.

Although this experiment was conducted before the RULEX model was first developed, its findings challenge RULEX's assumption that people learn categories by learning a rule and its exceptions even when not told to do so. Although subjects were given a substantial head start toward the solution of the problem in the rules-plus exceptions condition, that condition did not lead to much more rapid problem solutions. Also the rules-plus-exceptions strategy led to longer recognition latency times than the standard-instructions condition.

EPAM's simulations of all three conditions in this experiment were run with 1000 simulated subjects and the following free parameters: (1) speed-of-learning-parameter = 40, (2) weight for eye-height = 4, (3) weight for eye-spacing = 1, (4) weight for nose-length = 0.5, and (5) weight for mouth-height = 3.5.

This data has already been simulated with an earlier version of EPAM (EPAM V; Gobet, Richman, Staszewski & Simon, 1997) with good fits with the human data in both regards. Table 12 shows the overall fit of the classification probabilities of EPAM V, EPAM VI and the human data.

Table 12

Observed and Predicted Proportions of Correct Categorization for Each Face during Transfer

Face Number	Instruction									
	Standard			Rules & Exceptions			Prototype			
	People	EPAM5	EPAM6	People	EPAM5	EPAM6	People	EPAM5	EPAM6	
Old faces										
4A (1112)	0.97	0.83	0.86	0.89	0.87	0.87	0.77	0.75	0.85	
7A (1212)	0.97	0.96	0.89	0.94	0.98	0.90	0.97	0.85	0.88	
15A (1211)	0.92	0.96	0.93	0.94	0.94	0.87	0.88	0.90	0.90	
13A (1121)	0.81	0.78	0.67	0.72	0.64	0.62	0.70	0.68	0.74	
5A (2111)	0.72	0.76	0.66	0.78	0.77	0.74	0.60	0.68	0.68	
12B (1122)	0.67	0.71	0.65	0.73	0.80	0.76	0.45	0.57	0.43	
2B (2112)	0.72	0.68	0.66	0.70	0.67	0.54	0.72	0.57	0.53	
14B (2221)	0.97	0.89	0.79	0.91	0.94	0.82	0.83	0.75	0.67	
10B (2222)	0.95	0.96	0.89	0.95	0.96	0.86	0.87	0.89	0.72	
M	0.86	0.84	0.78	0.84	0.84	0.78	0.77	0.74	0.71	
New faces										
1A (1221)	0.72	0.54	0.60	0.45	0.26	0.45	0.73	0.56	0.70	
6A (1111)	0.98	0.93	0.93	0.88	0.74	0.82	0.87	0.88	0.90	
9A (2121)	0.27	0.52	0.43	0.08	0.33	0.36	0.28	0.50	0.46	
11A (2211)	0.39	0.58	0.55	0.75	0.88	0.74	0.52	0.61	0.59	
3B (1222)	0.44	0.63	0.58	0.80	0.91	0.75	0.35	0.55	0.41	
8B (2212)	0.77	0.55	0.67	0.42	0.28	0.45	0.78	0.53	0.52	
16B (2122)	0.91	0.83	0.81	0.88	0.66	0.73	0.88	0.74	0.67	

Note. Human data is from Medin and Smith (1981). EPAM V data is from Gobet et. al. (1997).

Table 13

Mean Reaction Times for Correct Responses for Each Old Face During Speeded Classification as a Function of Instructions

Face Values	Instruction								
	Standard			Rules & Exceptions			Prototype		
	People	EPAM5	EPAM6	People	EPAM5	EPAM6	People	EPAM5	EPAM6
4 (1112)	1.11	0.96	0.96	1.27	1.23	1.28	1.92	1.65	1.37
5 (1212)	1.34	1.02	1.20	1.61	1.33	1.55	2.13	1.68	1.60
7 (1211)	1.08	0.70	0.95	1.21	0.93	1.21	1.69	1.39	1.36
13 (1121)	1.27	1.03	1.20	1.87	1.44	1.51	2.12	1.64	1.53
15 (2111)	1.07	0.64	0.87	1.31	0.95	1.25	1.54	1.22	1.30
2 (1122)	1.30	1.14	1.22	1.97	1.53	1.59	1.91	1.78	1.68
10 (2112)	1.08	0.65	0.93	1.42	0.96	1.27	1.64	1.25	1.56
12 (2221)	1.13	1.10	1.24	1.58	1.30	1.38	2.29	1.75	1.63
14 (2222)	1.19	0.77	1.14	1.34	0.95	1.38	1.85	1.44	1.62
M	1.17	0.89	1.08	1.51	1.18	1.39	1.90	1.53	1.52

Note. Human data is from Medin and Smith (1981). EPAM V data is from Gobet et. al. (1997).

EPAM has a good fit with people’s latency times as shown in Table 13. Both people and EPAM recognize items more quickly in the standard condition as compared to the rules plus exceptions condition and the prototype condition. The progression for people is 1.17, 1.51, and 1.90 seconds. The progression for EPAM VI is 1.08, 1.39, and 1.52 seconds. This progression results from the way we have programmed the strategies.

The directions tell the subject to learn a rule for nose-length (the third dimension) and then memorize the exceptions. Our interpretation of this rule has EPAM create one discrimination subnet for the rule and

Table 14

Mean Number of Errors for Each Face during Initial Learning as a Function of Instructions

Face Number	Instruction								
	Standard			Rules & Exceptions			Prototype		
	People	Epam5	Epam6	People	Epam5	Epam6	People	EPAM5	EPAM6
4 (1112)	4.5	9.0	6.2	3.9	6.5	5.4	7.7	11.2	7.1
5 (1212)	8.2	10.8	9.8	5.9	8.5	7.1	9.2	11.9	10.4
7 (1211)	4.2	6.6	5.8	3.3	4.0	4.5	6.7	9.9	6.4
13 (1121)	11.9	11.3	9.1	10.7	18.8	9.5	13.7	12.8	9.2
15 (2111)	2.8	5.6	4.2	2.8	4.5	4.7	4.9	8.6	5.6
2 (1122)	12.9	14.0	10.5	13.8	18.5	11.6	10.3	16.0	15.2
10 (2112)	4.4	6.6	4.1	3.8	4.9	5.1	4.2	10.4	9.9
12 (2221)	15.2	12.8	11.4	6.3	7.7	6.1	17.4	15.1	17.0
14 (2222)	6.6	8.4	7.3	6.8	5.3	6.1	8.7	12.4	11.9
M	7.9	9.5	7.6	6.3	8.7	6.8	9.2	12.0	10.3

Note. Human data is from Medin and Smith (1981). EPAM V data is from Gobet et. al. (1997).

another discrimination subnet for exceptions. In the prototype condition, the instructions tell the subject to memorize what A faces look like and what B faces look like. Our interpretation of these instructions has EPAM create a discrimination net for A faces and another for B faces.

In order to recognize a stimulus in the standard condition, EPAM only needs to sort it through one discrimination net. However, in the other conditions EPAM has to sort the stimulus through two nets. The main reason for the difference in latency times is that sorting in two nets takes longer than the process of sorting in just one net. Not only that, but duplication of tests in the two nets requires additional eye fixations.

EPAM also has a good fit with errors made during learning as shown in Table 14. Like people, EPAM learns the rules-plus-exception stimuli with only a few less errors than the standard-condition stimuli. This might be surprising considering that the rules-plus-exceptions subjects are given the head-start of being told a good rule at the beginning of the experiment.

It would take RULEX quite a bit of time to find such a good dimension. First RULEX would have to iterate through all of the features in an attempt to find one that would be perfectly predictive. Next RULEX would have to again iterate through the features trying to find one that would be sufficiently predictive. Being given the dimension should give a rules-plus-exceptions strategy an excellent head-start.

EPAM does indeed take advantage of this head start to quickly learn the rule in a rules-subnet that it creates for this experiment. However, learning is not greatly speeded up because of the problems that EPAM encounters when learning to discriminate stimuli in the exceptions subnet.

Although EPAM uses the same strategy in the exceptions subnet that it uses in standard condition of this experiment,¹² there is a fundamental difference. In the standard condition subnets EPAM's task is to discriminate members of Category A from members of Category B, but in the exceptions subnet EPAM's task is to discriminate exceptions from non-exceptions. In the standard condition both Category A and Category B have central tendencies which assist EPAM's hypothesis-formation strategy. However in the exceptions subnet the two exceptions (stimuli 13 and 10) have nothing in common but the least predictive dimension (Dimension 2), and non-exceptions, consisting of all of the non-exception members of both Categories A and B have even less in common. It is much more difficult for EPAM to separate out exceptions from non-exceptions than it is for EPAM to separate Category A stimuli from Category B stimuli.

In addition there is one major difference between the learning strategy in the exception subnet from EPAM's strategy in the standard subnet: Rule 6 of the learning strategy. This rule only applies to those simulations in which the subject has been instructed to memorize exceptions to a rule. It is just like step 5 except that if the current stimulus and the previous stimulus are not members of the same category the hypothesis will be created by comparing the stimulus to itself. This allows the system to study the features of the exceptions and speeds up what would otherwise be very slow progress in the exceptions subnet.

EPAM is able to explain why a head-start for the rules-plus-exceptions condition does not reduce errors much and why rules-plus-exceptions instructions increase subject latency times. RULEX, however,

Table 15

Categorization Response Probabilities of reporting that item is a member of Category A during the Transfer Condition.

Stimulus	People	RULEX	Context	EPAM
Category A				
4 (1112)	0.94	0.97	0.94	0.96
7 (1212)	1.00	0.98	1.00	0.97
15 (1211)	0.97	0.99	1.00	0.98
13 (1121)	0.98	0.97	0.88	0.90
5 (2111)	0.92	0.93	0.87	0.84
Category B				
12 (1122)	0.13	0.13	0.18	0.18
2 (2112)	0.06	0.03	0.12	0.09
14 (2221)	0.02	0.01	0.00	0.06
10 (2222)	0.02	0.01	0.00	0.02
Transfer Item				
1 (1221)	0.94	0.93	0.94	0.79
3 (1222)	0.69	0.65	0.55	0.65
6 (1111)	0.94	0.98	0.97	0.90
8 (2212)	0.03	0.06	0.07	0.20
9 (2121)	0.14	0.26	0.42	0.41
11 (2211)	0.32	0.30	0.46	0.35
16 (2122)	0.08	0.02	0.06	0.23

Note. The Human data and the RULEX data are from Palmeri and Nosofsky (1995).

assumes that people would use a rules-plus-exceptions strategy in the standard condition even though they are not told to do so. Thus it has a difficult time explaining why the head start given to the rules-and-exceptions subjects doesn't help much and why the latency times for the rules-plus-exceptions-condition subjects are higher than those of the standard-condition subjects.

Although EPAM and RULEX share similar learning strategies that look for predictive features and conjunctions of two features, the difference in the way that they hold that learning (EPAM in a discrimination net and RULEX as a rule and exceptions) allows EPAM to simulate the learning times and latency times that result from these different experimental instructions.

The 5-4 Task and Old-New Recognition

The second weakness with RULEX was noticed by its authors themselves when Palmeri and Nosofsky (1995) found that they couldn't use RULEX to predict recognition data without merging it with Medin and Schaeffer's (1978) probabilistic-response-rule context model. The task that they explored in their first experiment was a rules-plus-exceptions

instructions task much like the task used by Medin and Smith with five modifications: (1) the stimuli were

Table 16

Median speeded categorization response times in Palmeri and Nosofsky's (1995) Experiment 1.

Stimulus	People	EPAM
Category A		
4 (1112)	1.47	1.03
7 (1212)	1.11	1.03
15 (1211)	1.28	1.03
13 (1121)	1.45	1.32
5 (2111)	2.12	1.32
Category B		
12 (1122)	2.24	1.32
2 (2112)	1.50	1.32
14 (2221)	1.14	1.07
10 (2222)	1.11	1.03
Transfer Item		
1 (1221)	1.26	1.28
3 (1222)	1.43	2.28
6 (1111)	1.32	1.03
8 (2212)	1.32	1.03
9 (2121)	1.36	2.06
11 (2211)	1.47	2.06
16 (2122)	1.55	1.03

Note. The Human data is estimated from Figure 2 of Palmeri & Nosofsky (1995).

drawings of rocket ships not Brunswick faces, (2) physical attributes were randomly assigned to logical attributes, (3) subjects were instructed to learn a rule for Dimension 1 instead of Dimension 3, (4) exceptions to the Dimension 1 rule were not presented during the first two training blocks, and (5) subjects were told that there would be exactly one exception on each side of the rule.

The probability of categorizing each face as an Category A during the transfer phase of the experiment is shown in Table 15 which shows the results for people, RULEX, and EPAM. All three models achieve good fits with the human data although RULEX's fit is better than either that of the context model or EPAM. EPAM's simulation of this experiment was run with 1000 simulated subjects and the dimensional weights equalized.

Palmeri and Nosofsky also reported the latency times for each stimulus during a speeded categorization phase of the experiment, but did not report the latency times predicted by RULEX. Latency results for people and EPAM are shown in Table 16. EPAM's fit with the latency times are quite poor in contrast to EPAM's good fits with latency data in the other simulations reported in this paper.

The striking feature of the observed latency times in this experiment was the lengthy period that the human subjects took to classify the two exceptions (Stimuli 5 and 12). Such long latency times were not observed by Medin and Smith (1981) in their rules-plus-exceptions condition shown in Table 13, where the two exceptions, Faces 13

Table 17

Old-new recognition probabilities observed and predicted
in Palmeri and Nosofsky's Experiment 1

Stimulus	People	RULEX	Context	EPAM
Category A				
4 (1112)	0.79	0.74	0.87	2.40
7 (1212)	0.70	0.65	0.72	2.34
15 (1211)	0.63	0.65	0.72	2.07
13 (1121)	0.85	0.75	0.83	2.57
5 (2111)	0.94	0.93	0.78	3.43
Category B				
12 (1122)	0.99	1.00	0.83	3.44
2 (2112)	0.75	0.74	0.80	2.65
14 (2221)	0.60	0.65	0.72	2.65
10 (2222)	0.79	0.65	0.63	2.32
Transfer Item				
1 (1221)	0.64	0.65	0.63	2.77
3 (1222)	0.72	0.72	0.63	2.69
6 (1111)	0.71	0.72	0.80	2.26
8 (2212)	0.56	0.65	0.63	2.04
9 (2121)	0.67	0.72	0.67	2.11
11 (2211)	0.66	0.70	0.63	2.06
16 (2122)	0.66	0.74	0.68	2.18
Exceptions	0.97	0.97	0.81	3.43
Confusables	0.80	0.75	0.82	2.61
Other Old	0.70	0.67	0.73	2.36
New	0.66	0.70	0.67	2.30

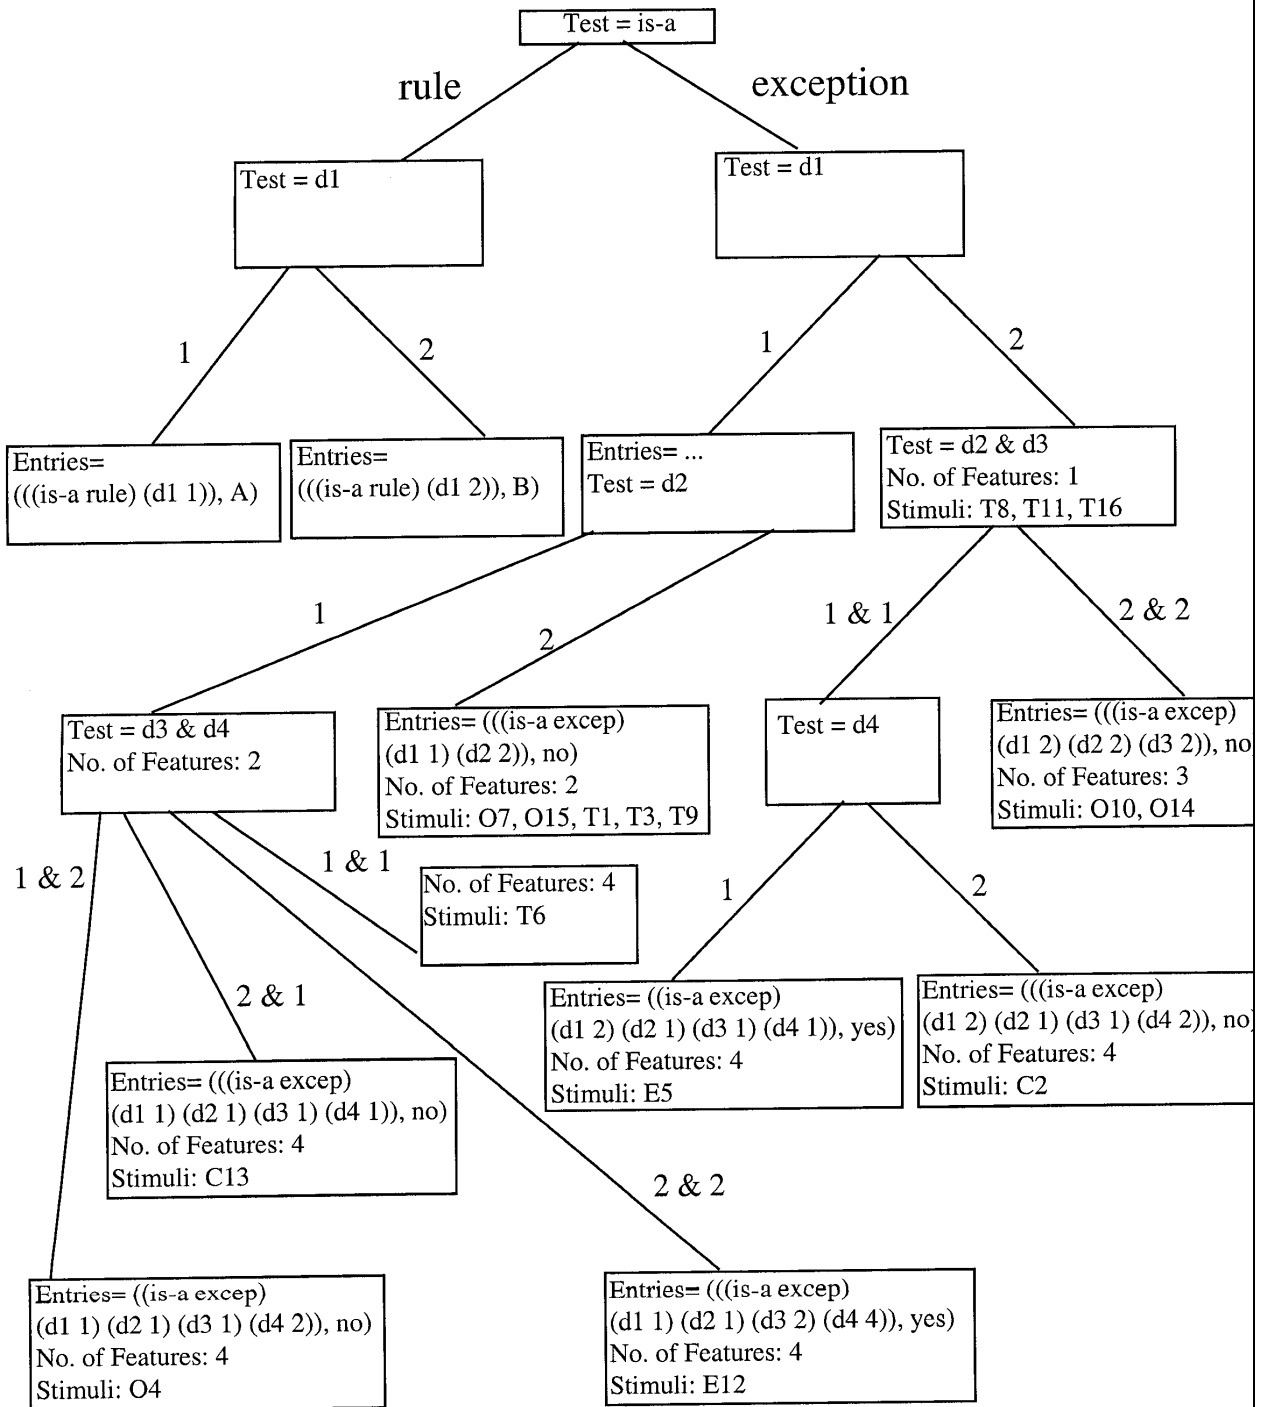
Note. The Human data is estimated from Figure 1 (A) of Palmeri and Nosofsky (1995).

and 10, were not classified more slowly than the other faces. Clearly something different is happening in this experiment than in the Medin and Smith experiment.

EPAM does not currently model why any of the differences between the experimental procedures of Palmeri and Nosofsky as compared to Medin and Smith would result in such a different pattern of latency results. These differences in procedure included: (1) the physical stimuli, (2) the dimension chosen for the rule, (3) randomization of assignment of physical to logical dimensions, (4) number of trials, (5) termination criteria, (6) criteria for including subjects in the results, (7) whether only rule-obeying stimuli were presented during the first two trials, (8) whether latencies are reported as medians or means, (9) whether a recognition trial was included before the latency collection trial, and (10) whether subjects were told how many exceptions there would be to the rule.

During the transfer phase, Palmeri and Nosofsky conducted a recognition test in which they asked subjects to report whether or not they had seen each of the 16 stimuli during the training phase. The

Figure 17. Discrimination net at the end of a run of Palmeri & Nosofsky's Experiment 1.



probability that an item would be identified as an “old” item is shown in Table 17.

The results in the EPAM column are in a different format from the results for people or the other models. Instead of reporting the probability of a “yes” decision about whether or not an item had been seen during training, we are reporting the raw data that EPAM would use to make such a decision. In the EPAM column we show the number of features that appear in what we call the “image” of the exception discrimination net node to which the stimulus sorts. They are precisely the features that must be present in order for a stimulus to sort to this node and are also the features that must appear on the stimulus side of the entry that appears at this node.

We presume that in order to make a decision about whether it had seen an item before, EPAM would count its familiar features. If all four features were familiar its score on this measure would be “4” and the item would clearly be one that EPAM had seen before. If no feature was familiar EPAM’s score would be “0” and the item would clearly be one that EPAM had not learned anything about. For an example of these calculations see the exception net from a run of this experiment pictured in Figure 17. Inside each node is the information about the number of features in the image of the node as well as information about which stimuli sort to this node (T8 indicates Transfer stimulus 8, E12 indicates Exception 12, etc.).¹³

For example, assuming EPAM is going to sort Stimulus 10. The values of the dimensions of this stimulus are 2, 2, 2, and 2. First of all EPAM will sort this stimulus in the rule subnet. The top test in this rule subnet is a test for Dimension 1. The stimulus’ value for this test is “2” so the stimulus is sorted to the node whose entry indicates that the rule for this stimulus is Category B. Next the stimulus is sorted through the exception subnet in order to determine whether or not the stimulus is an exception to the rule. The top test in this subnet is for Dimension 1 and so the stimulus follows the branch for the value “2” to a node whose test is for a conjunction of Dimension 2 and Dimension 3. The stimulus’ value for this test is 2&2, so the branch labeled “2&2” is followed to a node for items whose values for Dimensions 1, 2, and 3 are all “2.” At this node is an entry which says “no” — that this stimulus is not an exception. EPAM then combines the information that the item sorts to Category B in the rules net and that it sorts to “no” in the exception net and responds that the stimulus is a member of Category B. It’s confidence that it has seen this stimulus is “3” since there were 3 features at the exception net node to which this stimulus sorted.

The collected data appears at the bottom of Table 17. These rows average the results for four

categories of stimuli *exceptions* (stimuli 5 and 12), *confusables* (stimuli 2 and 13), *other old items* (stimuli 4, 7, 15, 14, and 10), and the new *transfer items* (stimuli 1, 3, 6, 8, 9, 11, and 16). The *exceptions* are those items that do not fit the rule for Dimension 1. The *confusables* are old items that only differ by one feature from *exceptions*¹⁴ and are thus the most difficult old stimuli to discriminate from the *exceptions*.

Table 17 shows that only EPAM predicts the same ordering as people where *exceptions* are most easily recognized as being old, *confusables* are second most easily recognized, *other old items* are third and *new items* are less often identified as old than are any of the categories of old items. Neither of the other models predicts this ordering. RULEX is low for *other old items* while the context model is low for *exceptions*.

In order to remedy the deficiencies in this recognition task of both the RULEX and context models, Palmeri and Nosofsky combined the two models. EPAM presents an alternative way of overcoming the deficiency of the RULEX model. Instead of merging RULEX with the context model, RULEX could be merged with a discrimination net model. In fact, that is exactly what we have tried to do in this study. We have merged the essence of RULEX's rule-finding strategy into EPAM.

Discussion

Strategies have long functioned as intervening variables in EPAM simulations. For example, an anchor-point strategy allowed EPAM I to explain the serial position effect in the serial anticipation paradigm (Feigenbaum & Simon, 1962) and a choice of rehearsal strategies enabled EPAM III to explain whether incremental learning would occur in a paired-associate experiment (Gregg & Simon, 1967). In the simulations reported here, the combination of a means-ends-analysis inspired strategy with discrimination net allowed EPAM to explain several classification learning experiments:

1. In our simulation of Nosofsky, Gluck et. al's (1994) replication of Shepard, Hovland, and Jenkins (1961) six tasks, EPAM provided a fairly good fit with the human data except that EPAM made two many errors overall on the sixth task. EPAM could perhaps fit this task better if it were able to switch to a rote-learning strategy when its standard hypothesis-testing strategy was repeatedly failing.
2. In our simulation of Martin and Caramazza's (1980) defining features and family resemblance experiments, EPAM was able to closely approximate the human data because it produced the same discrimination nets that had been deduced by Martin and Caramazza. In the defining features experiment, EPAM produced the same nets as human subjects but it did not put the "hair" attribute into the top discrimination net node as

often as people. EPAM could perhaps fit this task better if it would notice the similarity between the category name “Harry” and the attribute name “hair.” On the other hand, the fit was so close in the family resemblance simulation as to strongly support the main idea behind EPAM’s strategy, that people look for features and conjunctions of features that reduce the number of unclassifiable stimuli. Unlike EPAM, RULEX would not be able to reproduce the discrimination nets that Martin and Caramazza deduced from the family resemblance data.

3. In our simulation of Nosofsky, Palmeri, and McKinley’s (1994) experiment EPAM captured some of the generalization frequencies and the hamming distance distribution apparent in the human data. EPAM, however, like RULEX, failed to capture the frequency of the ABABBAB generalization, a generalization that can occur when subjects use a majority rule to classify the stimuli on the basis of the three most frequent features. EPAM also over-reported the AAAAAAA generalization, a generalization that is made when subjects respond with Category A to every stimulus. Although EPAM may be successfully capturing how some subjects do the 5-4 task, it certainly is not capturing how all subjects do the task.

4. In our simulation of Medin and Smith’s (1981) experiment which examined the effects of different strategy instructions upon task performance, we found the same results found by an earlier version of EPAM (Gobet et. al, 1997) and were again able to explain why rules and exceptions instructions led to longer response latency times and why they do not speed learning by very much even when they give the subject the head start of being told which dimension to use as the rule. While these results are explained easily by EPAM, they cast doubt upon RULEX’s assumption that people use a rules-plus-exceptions strategy even when not instructed to do so.

5. And finally, in our simulation of Palmeri and Nosofsky’s (1995) first experiment we found that EPAM, like the context model, does not closely fit much of the data. However, unlike either the context model or RULEX by itself, EPAM was able to explain the overall ordering of the old-new recognition data. EPAM is able to simulate this data because of the structure of the discrimination nets that it builds over the course of this experiment. Palmeri and Nosofsky proposed combining the deterministic RULEX model with the probabilistic context model in order to simulate this data. EPAM does not need to be combined with any other model to simulate this data since its discrimination nets naturally accommodate “old-new” discriminations.

EPAM’s strategy in these various simulations was based on the assumption that people approach a

classification learning task as a problem solving task and that they set themselves the goal of reducing the number of unknown stimuli. We then proposed two heuristics which, based upon our analysis of the Gluck et. al. protocols appeared to be used by at least some human subjects.

The Martin and Caramazza family-resemblance experiment is especially strong evidence of the first of our heuristics — that people look for a feature or a conjunction of two features that, when present, correctly classify some unclassified stimuli without misclassifying any of the yet unclassified stimuli. Although we discovered this heuristic from the discrimination nets produced by human subjects in the Gluck et. al. (2001) study, we could have just as easily discovered it by analyzing the discrimination nets inferred by Martin and Caramazza. In Figures 15 and 16, after the test for face shape has been placed in the discrimination net by human subjects, each of the features adopted into the discrimination net classifies some additional stimuli correctly. For example, if the face is fat, in Figure 15 (i.e. if the left branch is followed from the top test) then the “no hair” feature separates out Harry 5, 7, 8, 9, 11, and 12. Then a test for “smile” separates out Harry 1, 2, and 10. Except for the test for face shape at the top of the net, every test added to the discrimination net correctly classifies at least some of the unknown stimuli. Moreover, EPAM’s estimate of 250 ms per test that requires an eye movement, continued from EPAM V (Gobet et al., 1997), approximately fits the actual latency times of Martin and Caramazza’s subjects as shown in Tables 7 and 8.

The strategy that we have used in these simulations is just a first approximation. Further studies should explore the actual heuristics used by individual subjects in depth. It very well may be that different subjects use different heuristics.

The strategy that we have developed is limited in applicability to those experiments where the stimuli are presented incrementally, categories are artificial, training is supervised, categories are binary, and values of attributes are binary. In order to produce a more general strategy, EPAM needs to be fully combined with a problem-solver such as GPS (Newell and Simon, 1972), SOAR (Newell, 1990) or ACT-R (Anderson, 1993), which is able to apply a variety of heuristics as it searches to reduce the number of unknown items when doing a classification learning task.

In some ways the strategy used by EPAM is closely related to the search strategy used by Nosofsky, Palmeri, and McKinley’s (1994) RULEX model to find a rule. Both strategies search for predictive features and predictive conjunctions of two features. The main difference is that RULEX uses its strategy to only find

one test (which it calls a “rule”), equivalent to the top test in a discrimination net, while EPAM uses its strategy to build all of the tests in a discrimination net.

RULEX often achieves closer fits to experimental data than does EPAM, but this is partly due to the fact that RULEX has many more free parameters than does EPAM. For example, RULEX is able to closely fit the Shepard, Hovland and Jenkins data by changing its “branch” parameter from its default setting and is able to closely fit the hamming distance data by changing another parameter from its default setting. EPAM could fit data more closely if it were given more free parameters. Although other models of classification learning have achieved closer fits to some of these experiments, they generally have many more free parameters, sometimes so many free parameters that it is not clear whether it is their many free parameters or their theory of human processes that is fitting the data (Smith & Minda, 2000; Gluck et al., 2001). We have not used statistical tests to compare the variance explained by these models because such tests are not valid when models have differing numbers of parameters (see Richman & Simon, 1989, for a discussion).

RULEX was clearly a great advance. Although rooted in the early criterial-characteristic category models of the field, it could still explain the new family-resemblance-category tasks. In this paper we build upon its accomplishments. We have shown that when we add a RULEX-like strategy to EPAM, EPAM can simulate much of what RULEX was earlier able to simulate.

But because EPAM stores its learning in a discrimination net and uses a means-ends-analysis inspired strategy, we show that it may be able to take our understanding of human classification learning processes just a bit further down the road. EPAM is able to explain four results that RULEX would have difficulty explaining: (1) the discrimination nets created by Martin and Caramazza’s family resemblance experiment subjects and the response latency times that result from those nets, (2) the longer-latency times under rules-plus-exceptions instructions of Medin and Smith’s subjects, (3) the failure of the head start in the rules-plus-exceptions condition to help Medin and Smith’s subjects very much, and (4) the higher recognition ratings for “old” non-exception stimuli compared to “new” stimuli of Palmeri & Nosofsky’s first experiment subjects.

Also, RULEX has only been used to predict subject responses, not reaction latencies. However, in both the Martin and Caramazza (1980) simulation and the Medin and Smith (1981) simulation, EPAM’s discrimination net produces good approximations of people’s latency times through the assumption, supported by chess research and reading research, that an eye movement requires approximately 250 ms.

Rosch and Mervis’s (1975) demonstration that people could learn family-resemblance classification

structures led to the development of probabilistic response rule models, including the very successful context model of Medin and Schaffer (1978). No longer did a category need to have the defining cues explored by such pioneers as Bruner, Goodnow, and Austin (1956). However, first RULEX, and now EPAM VI, have demonstrated that defining cue models can, indeed, simulate family resemblance experiments if either the defining cues can have exceptions (RULEX) or the defining cues can be used as tests in a discrimination net

(EPAM VI).

References

- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. Cognitive Psychology, 6, 451-474.
- Anderson, J.R. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.
- Anderson, J.R. (1993). Rules of the mind. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Baddeley, A. D. (1983). Working memory. Philosophical Transactions of the Royal Society London, B302, 311-324.
- Baddeley, A. D. (1986). Working memory. Oxford: Oxford University Press.
- Baddeley, A. D. (1990). Human memory: Theory and practice. Boston: Allyn and Bacon.
- Bruner, J., Goodnow, J., and Austin, G. (1956). A study of thinking. New York: Wiley and Sons.
- De Groot, A.D., & Gobet, F. (1996). Perception and memory in chess: Heuristics of the professional eye. Assen: Van Gorcum.
- Feigenbaum, E. A. (1959). An information processing theory of verbal learning (Report No. P-I 857). Santa Monica, CA: RAND Corporation.
- Feigenbaum, E.A., & Simon, H.A. (1962). A theory of the serial position effect. British Journal of Psychology, 53, 307- 320.
- Feibenbaum, E.A., & Simon, H.A. (1984). EPAM-like Models of Recognition and Learning. Cognitive Science, 8, 305-336.
- Gluck, K.A., Staszewski, J.J., Richman, H., Simon, H.A., & Delahanty, P. (2001). The right tool for the job: Information processing analysis in categorization. Proceedings of the 23rd Annual Meeting of the Cognitive Science Society. London: Erlbaum.
- Gobet, F., Richman, H.B., Staszewski, J.J., & Simon, H.A. (1997). Goals, representations and strategies in a concept attainment task: The EPAM Model. In D.L. Medin (Ed.), The psychology of learning and motivation: Vol. 37 (pp. 265-290). San Diego, CA: Academic Press.

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, 5, 123-162.

Hunt, E.B., Marin, J., & Stone, P.J. (1966). Experiments in induction. Academic Press, New York, NY.

Just, M.A., & Carpenter, P.A. (1987). The psychology of reading and language comprehension. Boston: Allyn and Bacon.

Martin, R.C., & Caramazza, A. (1980). Classification in well-defined and ill-defined categories: Evidence for common processing strategies. Journal of Experimental Psychology, 109, 320-353.

Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning, Psychological Review, 85, 207-238.

Medin, D.L., & Smith, E.E. (1981). Strategies and classification learning. Journal of Experimental Psychology: Human Learning and Memory, 7, 241-253.

Newell, A. (1990). Unified theories of cognition. Cambridge, MA: Harvard University Press.

Newell, A., & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.

Nosofsky, R.M., Gluck, M.A., Palmeri, T.J., McKinley, S.C., & Glauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). Memory & Cognition, 1994, 22, 352-369.

Nosofsky, R. M., Palmeri, T.J., & McKinley, S.C. (1994) Rules-plus-exceptions model of classification learning. Psychological Review, 101, 53-79.

Palmeri, T.J., & Nosofsky, R.M. (1995). Recognition memory for exceptions to the category rule. Journal of Experimental Psychology: Learning, Memory, & Cognition, 21, 548-568.

Richman, H. B. (1991). Discrimination net models of concept formation. In D. H. Fisher, Jr., M. J. Pazzani, and P. Langley (Eds.), Concept formation: Knowledge and experience in unsupervised learning, San Mateo, CA: Morgan Kaufmann.

Richman, H.B., & Simon, H.A. (1989). Context effects in letter perception: Comparison of two

theories. Psychological Review, 96, 417-432.

Richman, H.B., Staszewski, J.J., & Simon, H.A. (1995). Simulation of expert memory using EPAM IV. Psychological Review, 102, 305-330.

Rosch, E., & Mervis, C. (1975). Family resemblances: Studies in the internal structure of categories. Cognitive Psychology, 7, 573-605.

Shepard, R.N., Hovland, C.J., & Jenkins, H.M. (1961). Learning and memorization of classifications. Psychological Monographs, 75, 113, Whole No. 517.

Smith, J.D., & Minda, J.P. (2000). Thirty categorization results in search of a model. Journal of Experimental Psychology: Learning Memory and Cognition, 26, 3-27.

Sperling, G. (1960). The information available in brief visual presentations. Psychological Monographs: General and Applied, 74, 1-29.

Sperling, G. (1963). A model for visual memory tasks. Human Factors, 5, 19-31.

End notes

¹ We have guessed that the hypothesis is held in the visual imagery-store. It is possible, however, that it is held elsewhere, perhaps in long-term memory, the auditory imagery-store, or in a combination of long-term memory and the imagery-stores.

² We are using the same procedure as RULEX to order the features. The difference is that RULEX chooses one particular ordering for each simulation run while we randomly choose a new ordering (based upon the weights) whenever a new hypothesis is formed.

³ Such information could, perhaps, be gathered by visual sensory processes that note changes when consecutive stimuli and categorizations are presented in the same visual field with the same visual coordinates.

⁴ In these simulations, sorting through a compound test usually takes 500 ms. (two eye fixations) plus some additional 10's of ms. for recursively sorting the compound value in a general net used for associating compound values with symbols.

⁵ These stimuli each consist of a set of four attribute-value pairs, three of which (with the attributes "color," "shape," and "size") are features of the stimulus. The fourth, whose attribute is "is-a," indicates the subnet of the discrimination net to which the stimulus will be sorted. It's value "shj" indicates that this stimulus will sort to a newly created subnet for "Shepard-Hovland-Jenkins" stimuli.

⁶ While the actual stimuli used with human subjects each consisted of three features, the features were not the color, size, and shape pictured in this figure.

⁷ A stimulus is presented by being placed in the visual sensory store. In addition to the attribute values noted here, each stimulus has three features that are not used by EPAM's subject routine component, but are "visible" to EPAM's experimenter routine component. These features are "time" which indicates the maximum time will be visible, "group" which indicates whether it is a stimulus or response, and "case-number" which is used by the experimenter when computing the results of a simulation.

⁸ In this discrimination net, two of the nodes point to the same entry in memory. This occurred because, when studying the association between large & white with Category A, EPAM "borrowed" the already present entry for Category A that was at the node for white. In EPAM VI an entry in declarative memory can be pointed to

by more than one discrimination net node.

⁹ The difference of latency times is not precisely 250 ms. because there is a subnet, not shown in this figure, for sorting the value of a compound test. In that subnet it takes 10 ms. per test, EPAM's default value for sorting through a test that does not require an additional eye fixation.

¹⁰ One of the tests shown in this figure has 2 tests: (1) a test for face-shape & nose and (2) a test for just face-shape. EPAM VI continues EPAM V's (Gobet et al., 1997) innovation of multiple tests at a node replacing the "nec" nodes of EPAM III (Feigenbaum & Simon, 1984). When there are multiple tests at an EPAM node, the tests are tried in order and the first test to have a branch for the value found on the stimulus is utilized.

¹¹ Kevin Gluck is currently working to prepare a technical report for publication which will include the protocol transcripts from this study.

¹² The standard condition uses the strategy that was used in the other simulations of this paper.

¹³ For simplicity we have used the number of features at a node as our indication of whether an item is "old" or "new". A slightly more complex routine would combine number of features with whether or not the stimulus sorts to a node that has an entry. In the net shown in Figure 17, four of the new stimuli sort to nodes that do not have entries (T6, T8, T11, and T16) while none of the old stimuli sort to such a node.

¹⁴ *Confusable* stimulus 13 (1121) only differs in Dimension 4 from *exception* stimulus 12 (1122); *confusable* stimulus 2 (2112) only differs in Dimension 4 from *exception* stimulus 5 (2111).