

Simulations of Paired Associate Learning using EPAM VI

Howard B. Richman, Herbert A. Simon*, and Edward A. Feigenbaum

Complex Information Processing
Working Paper #553

Department of Psychology
Carnegie Mellon University
Pittsburgh PA 15213

March 7, 2002

*Herbert A. Simon died on February 9, 2001, while this report was being prepared.

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

EPAM VI, a descendent of both SAL III (Hintzman, 1968) and EPAM III (Feigenbaum & Simon, 1984), explains 20 separate findings of the paired associate paradigm. Its components, each of which is justified by additional results explained, include a discrimination net, an overlearning process, chunking mechanisms, rehearsal strategies, declarative memory, exponential decay, slotted structures borrowed from EPAM IV (Richman, Staszewski, and Simon, 1995), and time tags theorized by Yntema & Trask (1963) and Lockhart (1969) that are estimated to have a standard deviation from accuracy of .97 in log time. Similar to SAM (Mensink & Raaijmakers, 1988), except that EPAM uses a discrimination net, not global matching, EPAM VI distills a theory from the half-century of paired associate research that can explain research of the past as well as such modern studies as Anderson and Reder's (1999) study of fan effects.

The work of a couple generations of verbal learning behavioral psychologists gathers dust on the shelves of our university libraries. From about 1925 until 1975 they taught to human subjects lists of paired nonsense syllables or words, a “stimulus” and a “response.” They tried out many experimental permutations and were able to deduce major principles of learning: “People have a harder time learning when the stimuli are similar,” they noted. “Sometimes learning one list causes items from a previous list to be forgotten” which they called “retroactive inhibition” – “or the next list to be more difficult” which they called “proactive inhibition.” Some of the terms that they coined are still with us today. But by and large their studies stand there, unread. What is there to be taken from their hard work? Is it possible to distill a theory from their efforts, a theory that can cast light on the problems that we investigate today?

During their era Bower’s (1961) “one element” model, a variation of Bush and Morteller’s (1959) “one trial learning” model, achieved some outstanding quantitative fits to human data with just one parameter. The basic idea was that people could form an association of a stimulus with a response in a single trial.

Right at the end of their era, a computer revolution was being born, and two doctoral students, Edward Feigenbaum at Carnegie Tech and Douglas Hintzman at Stanford were trying their hand at programming discrimination net models of the paired associate paradigm, and achieving quite a bit of success. The models that they developed, EPAM III (Simon & Feigenbaum, 1964; Feigenbaum & Simon, 1984) and SAL III (Hintzman, 1968) each simulated important aspects of that paradigm.

More recently, three global-matching models, SAM (Mensink & Raaijmakers, 1988), TODAM (Murdock (1982, 1997), and Minerva 2 (Hintzman, 1984), and two auto-associative neural network models (Chappell & Humphries, 1994; Rizzuto & Kahana, 2001) have been used to simulate paired associate learning. Although SAM, the Search of Associative Memory theory, has been used to simulate several paired associate experiments, none of these global matching or neural network models has yet simulated a sizable segment of this paradigm.

In this paper we combine the two successful discrimination net models, SAL III and EPAM III, add a few new features, and produce a new model, called EPAM VI, that can explain all that its predecessors could explain, and much more.

The resulting model will be more complex than either of its predecessors. Instead of just a few parameters, it will have many. We will use the same strategy that a mathematician might use to solve simulta-

neous equations. We will tackle one parameter at a time. Once we nail down a parameter value in a simple equation, we will plug that value into a slightly more complex equation making that equation tractable. So long as the number of equations exceeds the number of parameters, it is possible to solve a very complex set of equations in this fashion. In this paper, we will simulate 20 different experimental findings, 20 different equations.

The EPAM VI Model

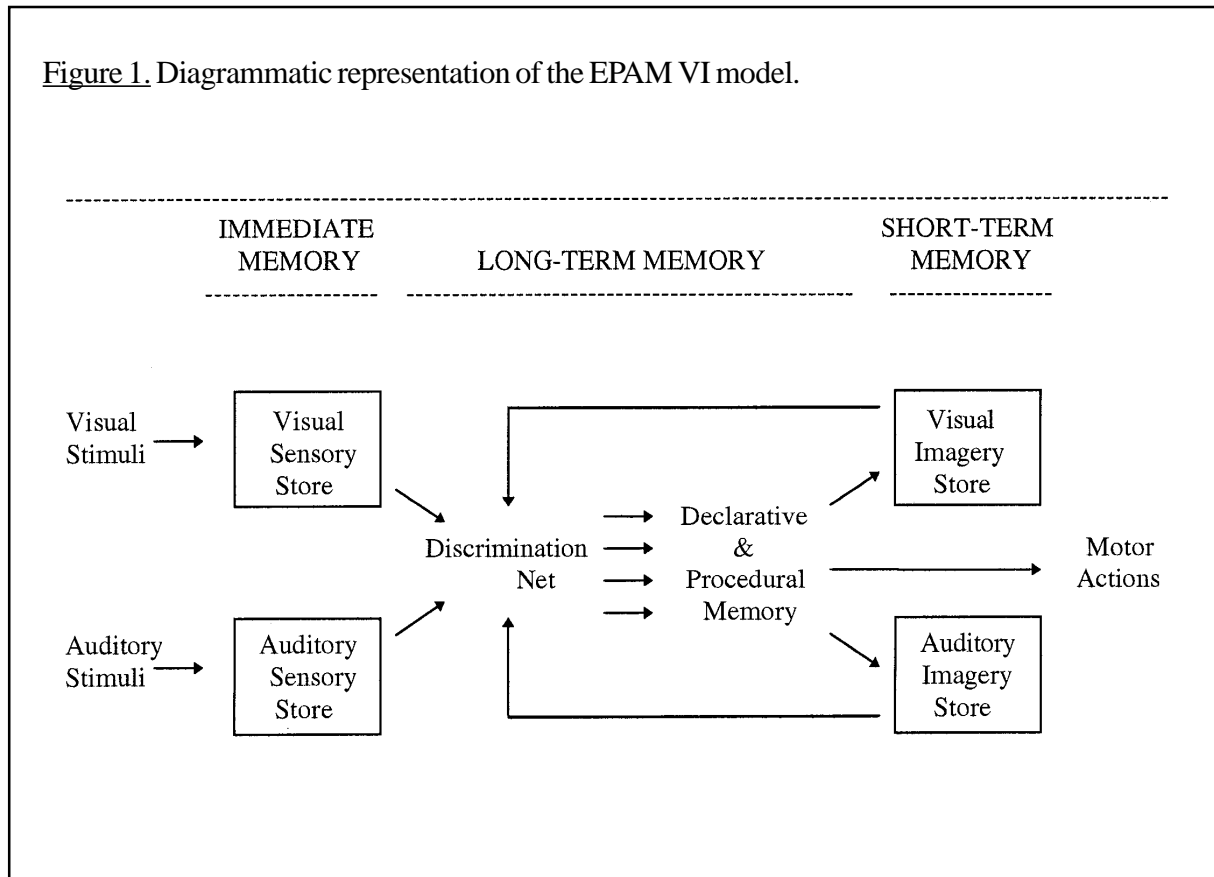
EPAM VI is the sixth revision of Elementary Perceiver And Memorizer, the first version of which was programmed by Feigenbaum (1959). Although EPAM VI models immediate memory, short-term memory, and long-term memory, its core is a part of long-term memory called the discrimination net which accesses the 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 classification learning, serial anticipation learning, distractor task, and articulatory-loop paradigms) are now in the process of being submitted for publication. The new revision 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. Like Atkinson & Shiffrin's (1968) two-stage model, there are both short-term and long-term memory stores. Short-term storage consists of sensory stores and imagery 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

Figure 1. Diagrammatic representation of the EPAM VI model.



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 was in the auditory imagery store is sorted through the discrimination net and then placed back, in refreshed form, 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 article, EPAM creates and/or expands a subnet for "syllables."

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 to simulate fan-effect experiments with the ACT models.

In the simulations reported here, EPAM VI creates entries in declarative memory each of which unites two discrimination net nodes, one for the stimulus and the other for the response. This entry would be indexed (i.e. pointed to) at both the discrimination net node for the stimulus as well as by the discrimination net node for the response.

Procedural Memory

EPAM VI 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 immediate memory stores and short-term memory stores in both the auditory and visual modalities. In the paired associate learning experiments simulated in this paper, the stimuli are usually presented visually as written letters (in the visual sensory store). The discrimination net then sorts the stimuli to entries of declarative memory where the response may be found and output. While each stimulus-response association is studied, the stimulus and response are subvocalized into the auditory imagery store and then rehearsed so that they are not lost from the store. These stimulus-response pairs take so little time to subvocalize that they can be rehearsed again and again without loss of information.

Plan for this Paper

In this paper we will introduce EPAM VI's components piece by piece in order to show why each of the components is needed. Whenever possible we will show when the same components enabled EPAM VI's more simple predecessors to explain the same findings.

We will begin with experiments that could be simulated by the simplest of discrimination net models, SAL I (Hintzman, 1968) and EPAM II (Feigenbaum, 1963). They establish that a simple discrimination net can explain the effect of intra-stimulus similarity upon error rates, the interaction of number of response alternatives with pre-criterion stationarity, various aspects of human protocols including stimulus-generalization, perseveration, and oscillation, and the effect of interstimulus similarity upon retroactive inhibition. We hold that these results establish a key component that we will plug into all of our other simulations – that people do indeed construct discrimination nets when they do a paired associate task.

Next we add an overlearning process that was incorporated into SAL II. When combined with the discrimination net, overlearning allows EPAM VI, like SAL II, to explain the pattern of total errors and intrusion errors that result from the degree of interpolated learning in a retroactive inhibition experiment.

The next experiments that we will simulate introduce components that were first worked out for EPAM III: that learning of a paired associate pair involves the learning of three chunks (stimulus, response, and association), that 5 to 10 seconds are required to learn a chunk, and that learning proceeds in incremental steps which are initiated during rehearsal. Using these additional parameters and their interactions with the discrimination net, EPAM VI will simulate Bugelski's (1962) finding that total learning time does not vary depending upon presentation rate, Chenzoff's (1962) finding that pre-familiarization of stimuli does not speed learning as much as prefamiliarization of responses, Underwood's (1953) finding that intra-stimulus similarity of stimuli slows learning while intra-stimulus similarity of responses does not, and Gregg and Simon's (1967) finding that differences in rehearsal strategy determine whether learning can occur in one trial.

Then we will add in Anderson's (1974) declarative memory with the idea, first included in SAL III, that a discrimination net node can point to more than one entry. With these additions we will be able to explain the interaction between backward recall and intra-stimulus similarity (Hintzman, 1969), differing estimates of recognition and recall, and the success of second and third guesses,

Then we solve a one parameter equation from the recency-discrimination research of Yntema

and Trask (1963) and Lockhart (1969) in order to derive a time tag parameter. When used in conjunction with our earlier components, time tags allow EPAM VI to explain the unlearning of old responses found by Barnes and Underwood (1959), the response latencies found by Suppes, Groen and Schlag-Rey (1966), the effects of inter-stimulus similarity upon transfer found by Bruce (1933), and fan-effects found by Anderson (1974) and Anderson and Reder (1999). Time tags, however, cause EPAM VI to undermatch high input probabilities in a probability matching experiment similar to one conducted by Voss, Thompson and Keegan (1959).

Next we add in a parameter hypothesized by Hintzman (1968) as part of the transition to SAL III, that newer responses at a node would be inherently unstable and decay exponentially. With this parameter, we explain the spontaneous recovery of old responses found by Briggs (1954).


Finally, we add in two components found in our study of expert memory using EPAM VI (Richman et al., 1995): (1) that experts can fill slots in just a quarter second, but (2) 21% of the information added to those slots is quickly forgotten. These two components allow us to explain paired-associate results when there are visualizable relations between stimuli and responses (Bower, 1972; Simon, 1972).

We will show that the sum of all of these components is able to explain 20 different experimental findings. All of the results presented in this paper were simulated using the final EPAM VI model with 1,000 simulated subjects unless otherwise reported. Some of the simulations reported in this paper make use of our full model of short-term memory¹ but most are “quick” simulations that do not simulate the entire system, just the discrimination net and declarative memory. One of the simulations (the *intralist similarity for stimuli and responses* simulation) was conducted in both forms – the normal and the quick – so that we could compare the results of the two versions. Both versions of the model use the same key routines for responding and learning. Those routines are described in detail in the appendix.

Regularities Simulated by Simple Discrimination Nets

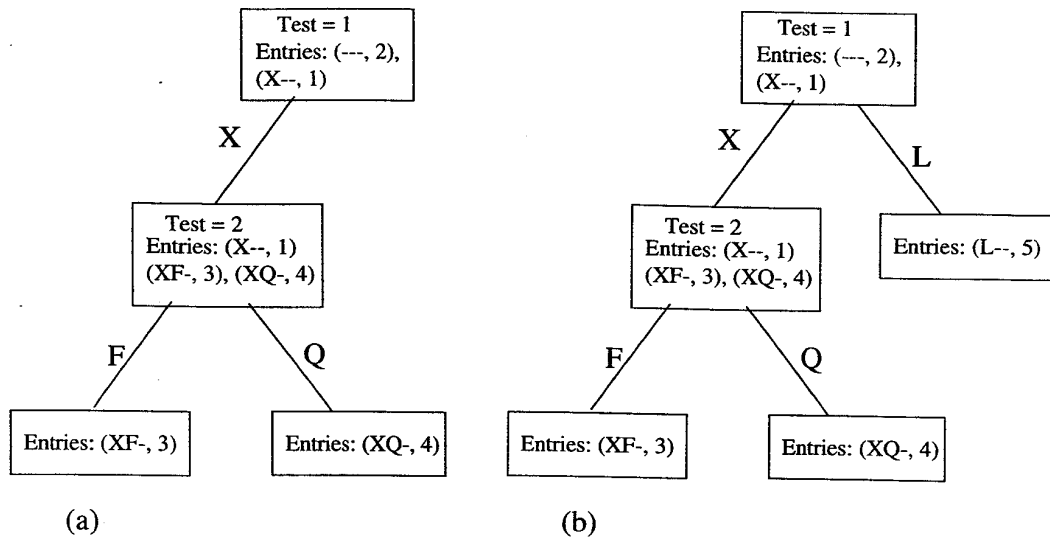
We begin with regularities simulated by the most simple discrimination net model, SAL I, the first of the SAL models developed by Hintzman (1968). The discrimination nets used by all of the SAL models were binary nets similar to those that had been earlier used in EPAM II. Figure 2 shows a portion of a SAL discrimination net which is capable of discriminating nonsense consonant trigrams so that the subject can respond with the assigned digit. It is shown both before and after SAL I has learned to respond with a “5” to any trigram beginning with the letter “L.”

Figure 2. A SAL discrimination net (a) before and (b) after growing a new test node for stimuli beginning with “L.” (from Hintman, 1967, p. 9).



In the net shown in Figure 2a, the simulated subject will respond to any trigram beginning with “XF” with “3,” any trigram beginning with “XQ” with “4,” any other trigram beginning with “X” with “1,” and any

Figure 3. An EPAM discrimination net (a) before and (b) after growing a new node for stimuli beginning with “L.”



trigram beginning with a letter other than “X” with “2.” Figure 2b shows the same net after the subject has engaged in discrimination learning (i.e. growing a new node in the discrimination net) in order to learn to respond with “5” to any trigram beginning with “L.”

The discrimination nets of EPAM VI are n-ary instead of binary meaning that every test can have more than two answers. Figure 3 shows EPAM VI nets that discriminate the same trigrams as the binary nets of Figure 2. Figure 3a shows the discrimination net before, and Figure 3b shows the same net after EPAM VI has learned to respond with “5” to any trigram beginning with “L.”

In SAL I, the learning algorithm involves two parameters, “a” and “b” as shown in Figure 4. After a responding error (responding with the incorrect response to a stimulus trigram), the simulated subject would draw a first random real number from 0 to 1. If the random number were less than “a” — a parameter between 0 and 1 — then discrimination learning would occur creating a new node in the discrimination net. If the first random number was greater than “a” then another decision point would be reached and another random number would be drawn. If the second random number was less than “b” then an old response in the net would be replaced by a new response. If the second random number was greater than b, then there would be no change in the discrimination net.²

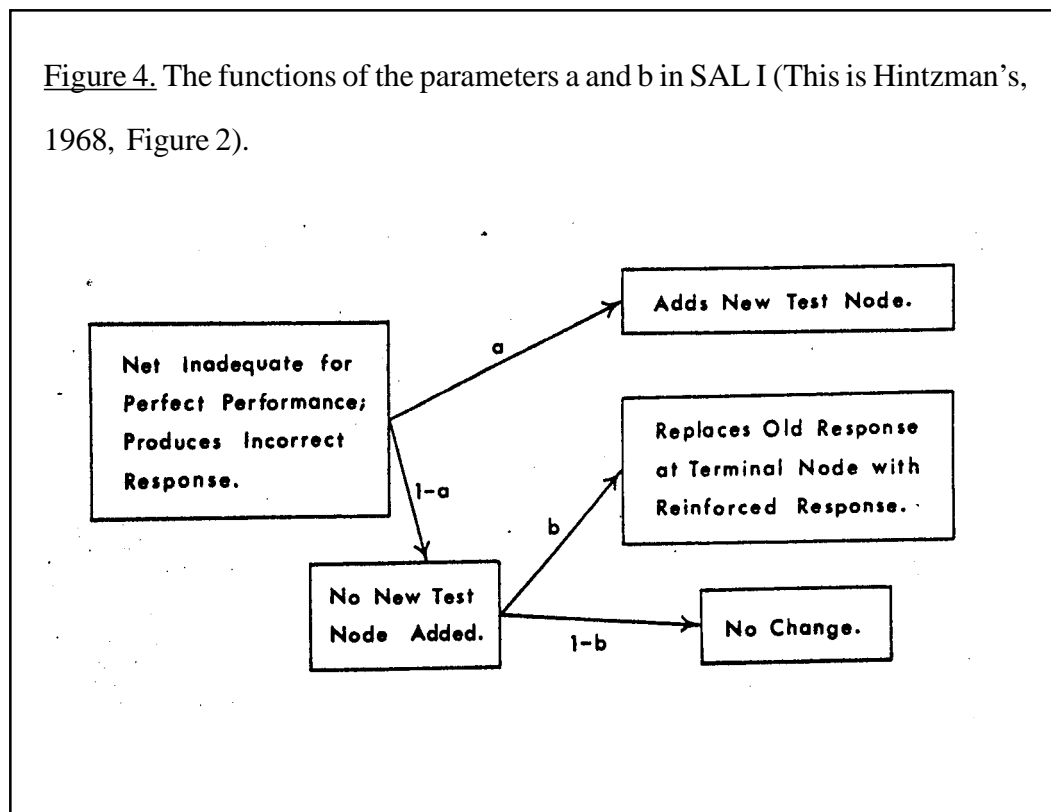


Table 1

Mean Errors per stimulus as a Function of Stimulus Similarity

Condition	People	EPAM VI
Low-similarity	4.84	4.68
High-similarity	9.19	7.42

Note. Results for people are from Hintzman (1969).

Hintzman used SAL I to successfully simulate intralist stimulus similarity and the interaction of precriterion curves with number of response alternatives. He also examined the protocols produced by SAL I and compared them qualitatively to those produced by human subjects in an experiment that he had conducted (Hintzman, 1967).³

Figure 5. EPAM VI discrimination net after run of low-similarity simulation

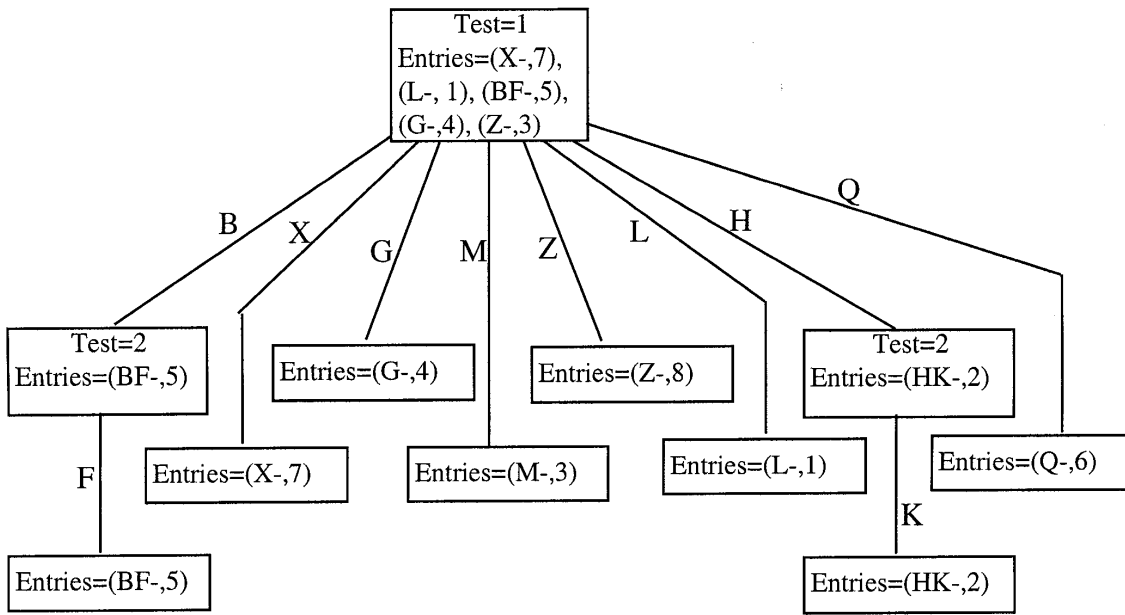
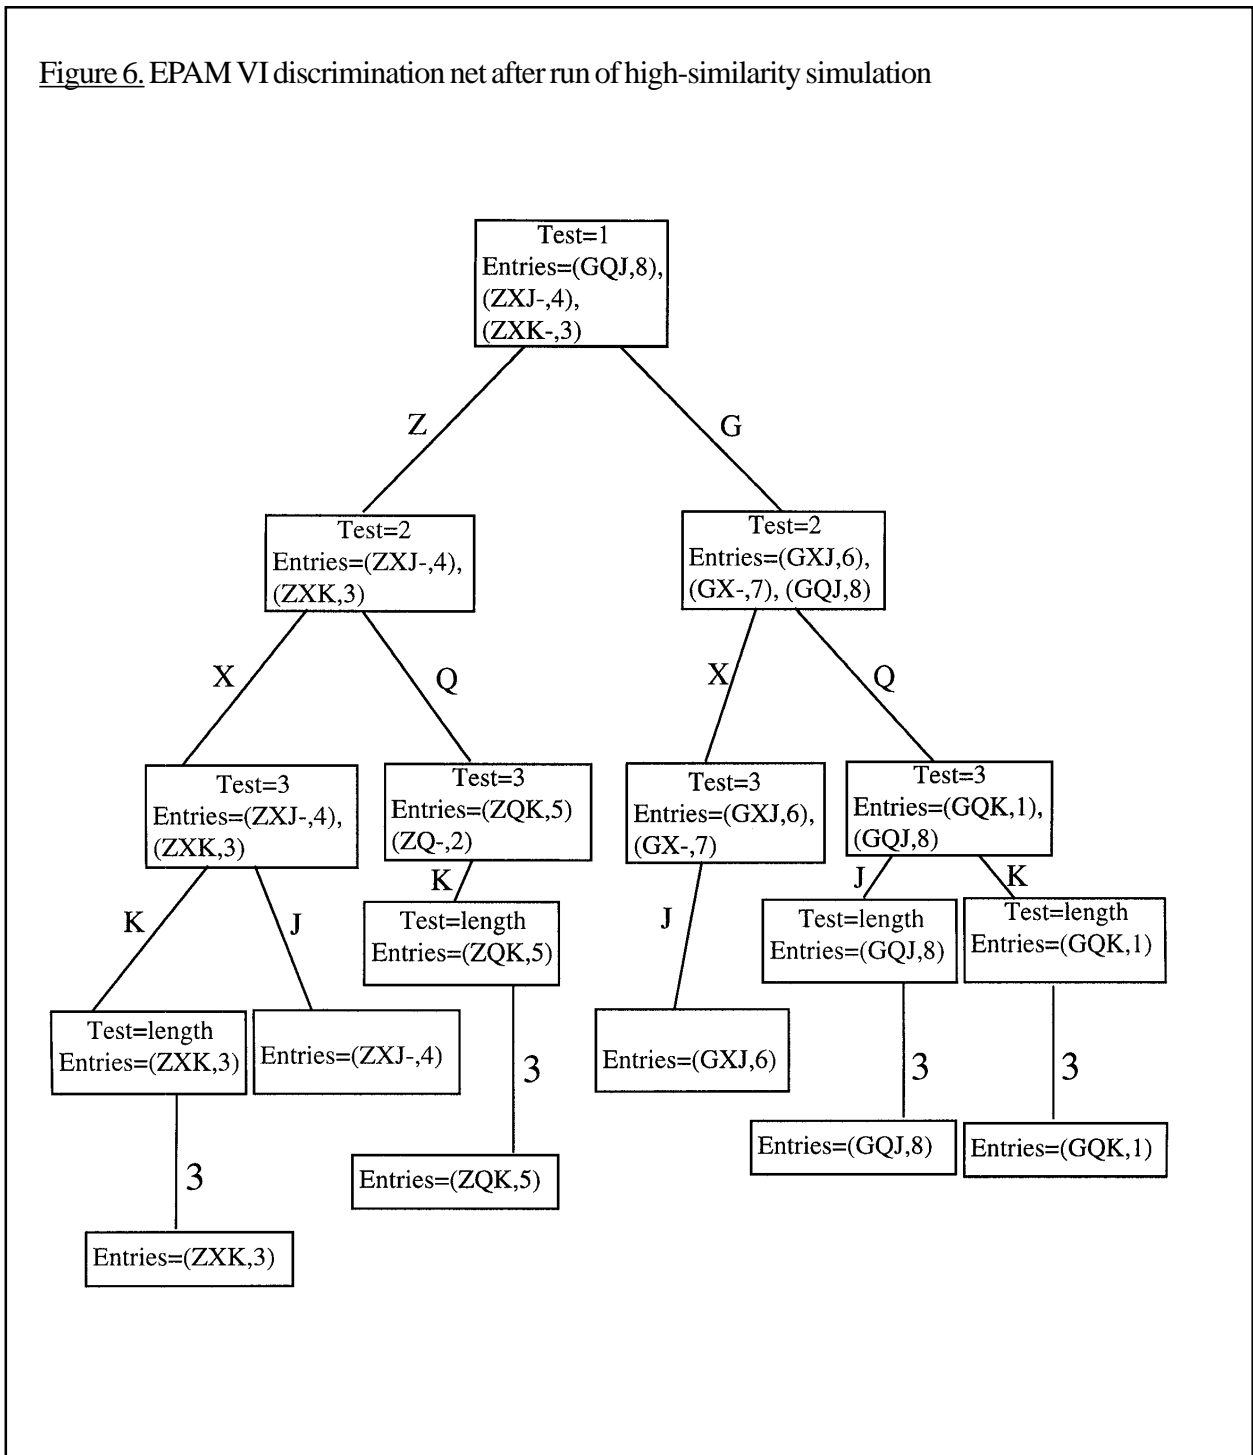


Figure 6. EPAM VI discrimination net after run of high-similarity simulation



Intralist Stimulus Similarity

When stimuli are similar to each other, people and EPAM make more errors during learning (Hintzman, 1969). Table 1 shows the results for EPAM VI and for people in an experiment that used nonsense consonant trigrams as stimuli and digits as responses. EPAM VI made an average of 4.68 errors per stimulus on the low-similarity list and an average of 7.42 errors per stimulus on the high similarity list. The averages for people

were 4.84 and 9.19. The stimuli used for in the low-similarity condition for EPAM VI consisted of eight trigrams created from different letters for every letter position. The stimuli for the high-similarity condition consisted of eight trigrams created from only two different alternative letters in each letter position.⁴

The difference between conditions occurred because the higher the similarity, the more test nodes were necessary in order to separate out the stimuli. Figure 5 shows a discrimination net created by EPAM VI during a low-similarity run while Figure 6 shows one created during a high-similarity run. A comparison of these two discrimination nets reveals that EPAM VI created 11 test nodes to discriminate the high-similarity stimuli from each other, but only 3 test nodes to discriminate the low-similarity stimuli.⁵

Table 2
A Sample SAL Protocol

Pair	Trial								
	2	3	4	5	6	7	8	9	10
GXJ-7	2	2	+	+	+	+	+	+	+
ZXK-8	4	6	6	+	6	+	6	+	+
GXK-4	2	+	+	+	+	+	+	+	+
LQF-5	9	+	+	+	+	+	+	+	+
GXF-2	1	+	+	+	+	+	+	+	+
LXJ-9	+	+	+	+	+	+	+	+	+
ZHJ-6	4	4	8	8	8	1	8	+	+
MBW-3	4	+	+	+	+	+	+	+	+
GQK-1	4	+	+	+	+	+	+	+	+
Total									
Errors	8	3	2	1	2	1	2	0	0

Note. This protocol is Hintzman's (1968) Table 1.

Qualitative Aspects of Protocols

When Hintzman (1968) compared the human protocols in the high similarity condition to those of SAL I (a sample SAL protocol is shown in Table 2), he found both similarities and differences.

Hintzman noted:

Several phenomena characteristic of human subjects can be seen here, including stimulus confusions or generalization (especially between ZXK and ZHJ), perseverative errors (e.g. repetition of the response 8 to stimulus ZHJ), and oscillation (e.g. the alternation of correct and incorrect responses to ZXK). These phenomena are also found in EPAM protocols (Feigenbaum, 1963). There are no failures to respond because SAL I does not deal with response latencies, and therefore assumes there is always enough time for the response to occur.

Generally, SAL protocols look much like human protocols, though some discrepancies can be detected. One discrepancy is that while human subjects will often reach a criterion of one errorless trial only to follow this with one or more errors, SAL I never does... Human subjects have difficulty learning to discriminate pairs of trigrams such as LQF and QFL since the same letters are used in both. SAL, on the other hand, would not confuse these two trigrams very frequently because its letter-matching tests are all position specific. (p. 129)

The EPAM VI protocol shown in Table 3 (for the simulation run whose final net is pictured in Figure 6) shares with SAL the stimulus-generalization, perseveration, and oscillation of the human protocols. In addition, it captures an element of the human protocols not captured by SAL I — the tendency for people to return to making some errors even after the first perfect trial. In the sample protocol shown in Table 3, EPAM had a perfect trial on Trial 11, and yet made two errors on Trial 12. This pattern of results sometimes happens when there is more than one entry at a node.

Specifically, during trials 11 and 12 all stimuli beginning with “ZQ” sorted to a node where there were two entries: (ZQ-, 2) and (ZQ-, 5). During trial 11, when “ZQJ” was presented, EPAM responded correctly with “2” because it perceived (ZQ-, 2) as being the more recent of the two entries and it also responded correctly to “ZQK” because at the time it was presented, EPAM perceived (ZQ-, 5) to be the more recent of the two entries. However, on the next trial, EPAM’s perceptions of the recency of both entries was the

Table 3

A Sample EPAM VI Protocol

Pair	Trial													
	2	3	4	5	6	7	8	9	10	11	12	13	14	
ZXK-3	8	4	+	4	+	+	+	+	4	+	+	+	+	
ZQJ-2	4	8	3	+	+	+	5	+	5	+	5	+	+	
GQK-1	3	6	7	7	8	+	8	+	+	+	+	+	+	
GXJ-6	7	7	7	7	+	7	7	7	+	+	+	+	+	
ZXJ-4	3	3	+	3	3	+	3	+	3	+	+	+	+	
GXK-7	4	+	+	+	6	6	+	+	+	+	+	+	+	
ZQK-5	3	4	2	2	2	2	+	+	2	+	2	+	+	
GQJ-8	7	6	7	7	+	1	+	+	+	+	+	+	+	
Total														
Errors	8	7	5	6	4	4	4	1	3	0	2	0	0	

opposite, and EPAM responded incorrectly to both stimuli. However, just after responding incorrectly to “ZQK” on trial 12, EPAM studied the correct response and created a new node in the discrimination net for “ZQK” which allowed it to respond correctly on subsequent trials.

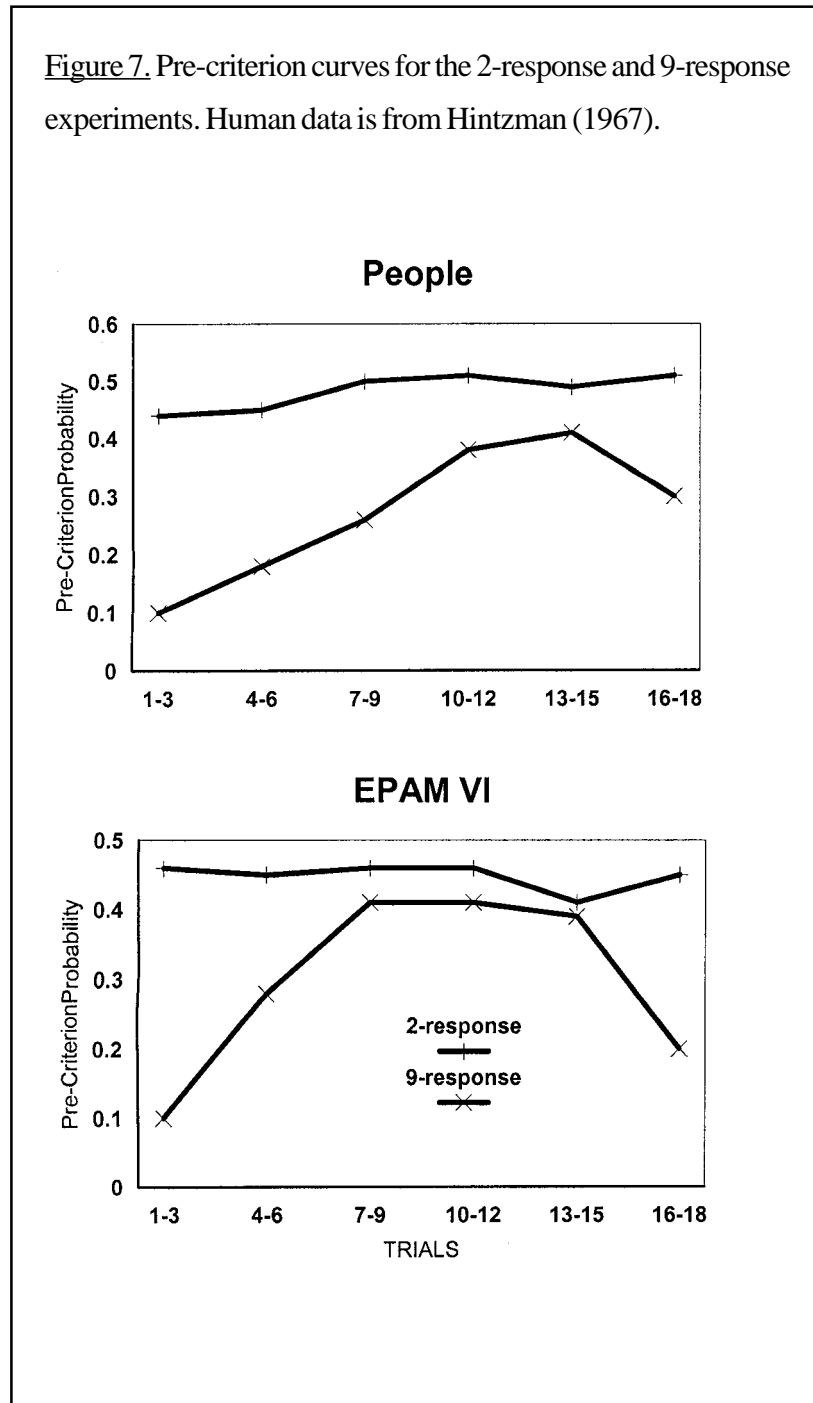
However, EPAM VI’s protocols do share one of the inadequacies pointed out by Hintzman for SAL I’s protocols: Both EPAM VI and SAL test for the occupants of letter-positions and thus fail to demonstrate confusion between pairs of trigrams like LQF and QFL which have the same letters but in different positions. A sister discrimination net model, WEPAM (Wynn,1966), could have confused LQF and QFL. It included not only tests for letter position but also tests for the presence of particular letters regardless of position.

Number of Response Alternatives

One of the predictions of Bower’s (1961) one-element model is called “pre-criterion stationarity” which means that, before the trial during which learning took place (i.e. “during the pre-criterion stage”)

subjects will simply *guess* a response resulting in a flat (i.e. “stationary”) curve in which the proportion correct is the probability of guessing correctly, but after the learning trial subjects will *know* the response resulting in always-correct responses.

The “one element model” was found only to apply when stimuli were highly dissimilar from each other. However, Hintzman (1967), following up a suggestion made by Bower (1967), conducted an experiment with human subjects which found that “pre-criterion stationarity” would also occur when there were just two re-



sponse alternatives, a phenomenon that he could explain using SAL I.

Figure 7 shows the pre-criterion curves for Hintzman's (1967) human subjects and for EPAM VI. The essential finding is that the pre-criterion curves for the 2 response lists of both people and EPAM VI hover at around .50 correct (the level of chance for a 2 response list) while the pre-criterion curves for the 9 response list quickly rise above the .11 level of chance for a 9 element list but remain below .50.

Both SAL I and EPAM VI easily produce these pre-criterion curves. They do not produce pre-criterion stationarity in the 9-response condition because growth of the discrimination net reduces the number of items with which each item is confused. For example, once a branch has been created in a discrimination net for trigrams beginning with "G," those trigrams beginning with "G" will all sort down that branch where they will only be confused with other trigrams beginning with "G."

Discrimination nets also predict that the top level that the pre-criterion curves can reach is .5. As Hintzman (1968) observed:

An error by SAL indicates that at least two stimuli are still being confused; therefore the probability of a correct response prior to the last error cannot exceed .50. This value is a theoretical upper limit for pre-criterion curves. Hence, as learning progresses, the 9-R curve rises, approaching .50, while the 2-R curve, which begins at .50 on early trials, remains stationary. (p. 134)

Interlist Stimulus Similarity and RI

One of the greatest thrills of scientific model builders occurs when a model that is designed to do one thing shows that it can explain other phenomena as well. Feigenbaum (1959, 1963) got that thrill when EPAM II was able to explain the forgetting that occurs when similar stimuli are learned. When that forgetting occurs while learning a single list it is called "oscillation." When learning of a second list causes first list responses to be lost it is called "retroactive inhibition." In both cases it is caused by the same mechanism: New learning in the same region of the discrimination net can make old learning inaccessible.

Retroactive inhibition (RI) was usually studied using the AB-CD-AB paradigm: First, the list AB was learned (stimuli A are associated with responses B). Then the interpolated list (CD) was learned (stimuli C are associated with responses D). Then list AB was relearned and the errors that occurred on the first trial were measured. These errors were seen to be a measure of the amount of forgetting that occurred with the AB list

as a result of learning the interpolated list.

All discrimination nets predict that retroactive inhibition will be higher, the higher the degree of interlist similarity between the A and C syllables as reported by McGeoch and Irion (1952). This forgetting occurs because similar syllables from List C would be learned in the same region of the discrimination net as the syllables from the List A. The results for SAL II and EPAM VI are shown in Table 4. In the low-similarity condition both models show very little interference, SAL II gets 99% correct upon relearning and EPAM VI gets 96% correct upon relearning. On the other hand, both models show a large effect when the lists are highly similar, SAL II gets 49% correct while EPAM VI gets 59% correct.


Condition	SAL II	EPAM IV
High-sim	0.49	0.59
Medium-sim	0.76	0.87
Low-similarity	0.99	0.96

Not all discrimination nets, however, would predict the pattern of total errors and intrusion errors that occurs depending upon how much the interpolated list is studied (Thune and Underwood, 1943). In order to explain that finding, Hintzman had to add an overlearning process to the SAL I model.

Overlearning and Retroactive Inhibition

Hintzman dubbed SAL II “The Overlearning Model.” The new feature of SAL II was a parameter labeled “c” which permitted learning (i.e. growing a new branch and node in the discrimination net) not only to occur after an incorrect response, but also after a correct response with a probability of “c.”

Figure 8. RI as a function of number of interpolated trials. Human data is from Thune and Underwood (1943). SAM data is from Mensink and Raaijmakers (1988).



We have incorporated the same overlearning step and parameter into EPAM VI. We call the parameter “*SAL-C*” and have set it at 30% in all EPAM simulations. That parameter determines the percentage of time that EPAM VI will study either the stimulus or the response (with a 50-50 chance of studying the stimulus or the response). In either case, the entry will be given a fresh time-tag. With this parameter in place, SAL II and EPAM VI were used to roughly simulate an experiment that Thune and Underwood (1943) conducted which revealed the effects of degree of learning of the interpolated list upon relearning.

Several aspects of Thune and Underwood’s experiment were modified for the simulations. While Thune and Underwood used the AB-AC-AB design, EPAM VI and SAL II used the AB-CD-AB design. While Thune and Underwood used real words as stimuli, EPAM and SAL used consonant trigrams such that each of the CD stimuli shared its first two consonants with an AB stimulus. Another difference was that Thune and Underwood only had their subjects study the original list for five trials, while SAL and EPAM studied the original list until a criterion of one perfect trial was reached. Yet another difference was that Thune and Underwood only counted an intrusion of the same response from the interpolated list once, even if the same response

popped up several times during the first trial of relearning list AB. In order to simplify our computations, we counted an intrusion each time it appeared.

As shown in Figure 8, the total number of errors for people (Thune and Underwood, 1943) and for EPAM VI rises and then stays relatively constant. However, the number of intrusions rises quickly with total errors and then gradually falls off as the interfering list is overlearned.

Figure 8 also shows the prediction of the SAM model with added contextual fluctuation processes (Mensink & Raaijmakers, 1988). These were not easy results for the SAM model to simulate. Its authors noted:

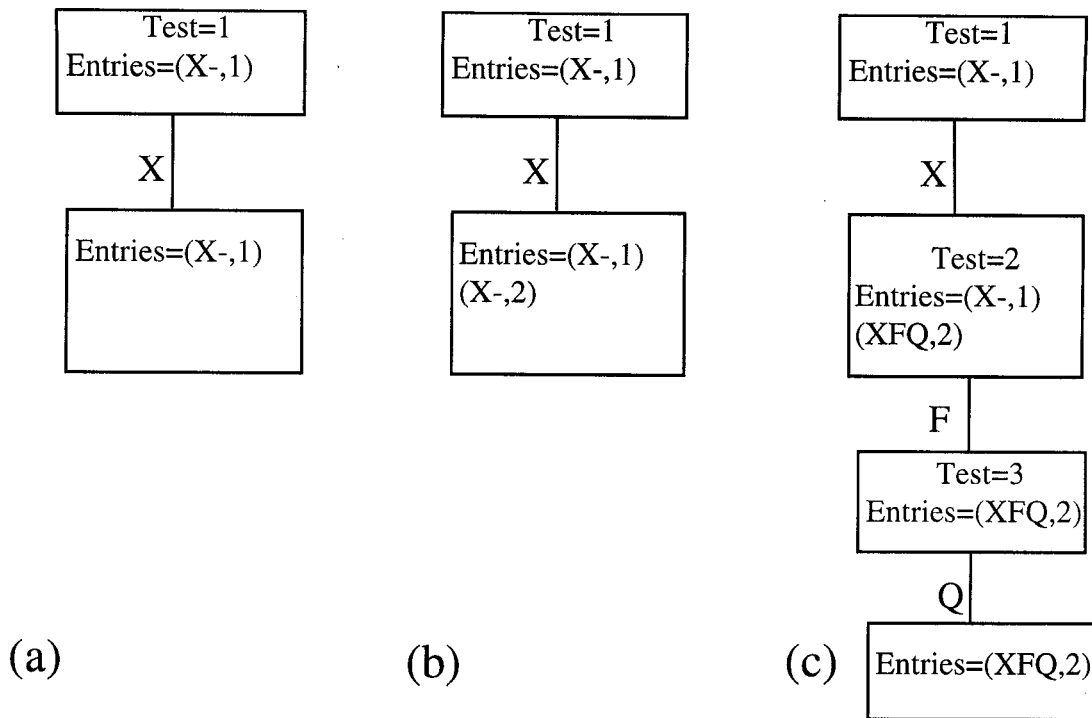
It must be mentioned that in our experience the shape of the intrusion function depends upon the parameter values used; that is, the inverted U-shaped function is not a parameter-free prediction. According to the model, a monotonically decreasing function may also be observed....(p. 441)

EPAM VI shows a good fit with the shape of these curves, although Thune and Underwood's subjects make fewer errors than EPAM VI. It is possible that this difference of scale is due to one or other of the many differences between the tasks given to Thune and Underwood's subjects and EPAM VI.

For EPAM VI, the initial rise in intrusions and overall rise in errors occur because the stimuli of the CD list sort to nodes where they compete with the responses of the AB list. The decline in intrusions occurs as a consequence of the over-learning illustrated in Figure 9. Figure 9a shows a portion of the EPAM net after $XFT \rightarrow 1$, a first-list stimulus-response pair has been learned. At this point, any stimulus beginning with X would be responded to with the response "1." Figure 9b shows the net after EPAM has begun to study $XFQ \rightarrow 2$, an interpolated-list stimulus-response pair. At this point, if the first list stimulus (XFT) were presented again, an intrusion error would occur if the entry (X-,2) is perceived by EPAM to have the more recent time tag. Figure 9c shows the same net after much overlearning of $XFQ \rightarrow 2$. At this point the original list stimulus (XFT) would sort to the node just above the bottom node (the one with "Test = 3"). Although there is indeed an "intrusive" entry at that test node, EPAM VI does not respond with the incorrect response because the stimulus part of (XFQ, 2) does not match the stimulus XFT, and so that entry is rejected. In such a case, EPAM VI would either not respond at all or would guess a response.

Both EPAM VI and SAL II were run using the AB-CD-AB design using nonsense syllables as stimuli and responses, not Thune and Underwood's AB-AC-AB design using familiar words as the stimuli and re-

Figure 9. Portion of an EPAM VI net showing overlearning of XFQ.



sponses. Thus both models, as currently constituted, share the same problem, described by Hintzman (1968):

It should be noted that the learning results of [this experiment] would not have been obtained if the A-B, A-C paradigm had been used – that is, if the stimuli of the two lists had been identical. Overlearning is only effective for SAL insofar as it increases differentiation between list-1 and list-2 stimuli, and SAL cannot discriminate stimuli which are formally identical.

To make the model applicable to overlearning in the A-B, A-C situation, one would have to assume that the model can add to a stimulus some kind of identifying tag or label which serves to differentiate the two lists. This identifying label could be incorporated into the discrimination net much as letter tests are. But list identification would ultimately have to be derived from temporal cues or from “context cues”—stimuli not obviously part of the immedi-

ate nominal stimulus situation (other pairs in the list, etc.). Indeed there is much evidence that a complete model of verbal learning would have to make use of context cues (e.g. Bilodeau and Schlosburg, 1951), but it is not clear how a mechanism to do this efficiently could be incorporated into the present model. (p. 144)

Effects of Chunking and Rehearsal Strategy

All versions of EPAM are chunking models based upon the idea that a node is a chunk of information. The following are among the basic findings that have all been found regarding chunks (for a review of some of these findings and others see Gobet, Lane, Croker, Cheng, Jones, Oliver & Pine [2001]):

1. It takes at least 5 to 10 seconds to memorize a chunk of information (Simon & Feigenbaum, 1964) unless that information can be added to the slots of a pre-existing chunk structure (Richman et al., 1995).
2. The amount of information that can be maintained in short-term auditory memory is a function of both the number of chunks and also the number of syllables (Zhang & Simon, 1985).
3. Only 2 to 3 chunks can be maintained in short-term visual memory (Zhang & Simon, 1985).
4. A maximum of two chunks can be perceived visually during one eye fixation (Richman & Simon, 1989).
5. The speed up of performance through practice can be explained as a result of the formation of new chunks of information (Rosenbloom & Newell, 1982; Newell, 1990).
6. The greater ability of a chess master, as compared to a chess novice, to reconstruct chess positions after seeing them for just a few seconds can be attributed to the greater number and greater size of chess-related chunks in the chess master's memory. (Chase & Simon, 1973; Gobet & Simon, 1998).

The paired-associate learning process of EPAM VI, like that of EPAM III, involves learning three chunks, a chunk for the stimulus, a chunk for the response, and a chunk for the entry. Each of these chunk learning processes takes a series of steps amounting to approximately 5 to 10 seconds of learning time. The time required for several steps of learning is the sum of the time required for each component step.

For example, in a paired associate simulation involving the associative pairing of the stimulus "BAJ" with the response "JID," EPAM VI will usually sort the stimulus and response by their first letters to nodes for

“B-” and “J-” and will create an entry (B-,JID) which unites the stimulus with the response and includes sufficient information to respond with the entire response. Specifically, the following steps would usually be required when learning to associate “BAJ” with “JID”:

1. 4 seconds to learn a node for (B-) plus an additional 8 seconds if a new test had to be added to a node.
2. 4 seconds to form the new entry in declarative memory ()
3. 4 seconds to add (B-) to a new entry resulting in (B-,)
4. 4 seconds to learn (J-) plus an additional 8 seconds if a new test had to be added to a node.
5. 4 seconds to add (J-) to the entry resulting in (B-,J-)
6. 4 seconds to add I to entry resulting in (B-,JI-)
7. 4 seconds to add D to entry resulting in (B-,JID)

Total learning time, here, to learn the three nodes was 28 seconds.⁶

Constant Learning Time

In an experiment that has not been previously simulated by EPAM but which all previous versions of EPAM could have simulated, Bugelski (1962) found that the total time required to learn a paired-associate list is not affected by presentation rate. He presented the same list of syllables to different human subjects varying only the presentation rate and found that the time required to learn the list remained constant. He was also able to extrapolate from this experiment in order to roughly estimate that the time required to learn each pair was approximately 25 seconds, which is very close to the 28 seconds, calculated in the paragraph just above for EPAM. Specifically he wrote:

The correspondence between total times for the most rapid presentation time and the longest presentation time is striking. Although the former required more than three times as many trials as the latter, the mean time difference is hardly more than 1 sec. Extrapolating roughly from these findings it might be possible for similar S's to learn all eight pairs in one trial with a presentation time of about 25 sec. per pair. (p. 410)

(Insert Table 5 about here)

As shown in Table 5, EPAM VI produces a very good fit to Bugelski's finding that time was constant, regardless of presentation time, to the total amount of time each pair was exposed during an experiment.

Table 5

Mean Trials and Total Times to Learn Lists

Presentation Time	People		EPAM VI	
	Trials	Exposure-Time	Trials	Exposure-Time
6 sec	10.2	61.2	9.3	55.6
8 sec	8.8	70.1	7.2	57.8
10 sec	5.8	57.9	5.9	58.6
12 sec	4.7	56.1	5.0	59.5
19 sec	3.3	62.2	3.5	65.9

Note. Results for people are from Bugelski (1962)

Bugelski's subjects ranged from 56.1 seconds to 70.1 seconds of total exposure per pair, while EPAM VI took from 55.6 seconds to 65.9 seconds.⁷

Of course not all subjects will show the same learning ability as the sophomore experimental-psychology course students from Buffalo University that Bugelski used in this study. Perhaps other population groups would be better, or worse, at attending to the task throughout the course of an experiment. Also there is evidence, which we will discuss later, that Bugelski's syllables suggested meanings to his subjects that could have made them easier to learn. It could well be that the learning time parameters that we have estimated for EPAM VI should be lengthened.

Prefamiliarization of Stimuli and Responses

One of the surprising results from paired-associate simulations was that advanced familiarity with stimuli was not nearly as helpful as advance familiarity with responses (Chenzoff, 1962). EPAM predicts this difference simply because responses must be learned completely (in order to respond correctly the entire response must be spoken) though stimuli only need to be learned sufficiently to be differentiated from one another.

If the response, JID, is already a completely familiar so that it is not just (J-) but (JID) then the learning time to learn the three chunks would be:

1. 4 seconds to learn (B-) plus an additional 8 seconds if a new test node is created.
2. 4 seconds to form entry ()
3. 4 seconds to add (B-) to entry resulting in (B-,)
4. 4 seconds to add (JID) to entry resulting in (B-,JID)

Total learning time here is 16 seconds.

If the stimulus, BAJ, is already completely familiar, the improvement in learning time is not as pronounced:

1. 4 seconds to form entry ()
2. 4 seconds to add (BAJ) to entry resulting in (BAJ,)
3. 4 seconds to learn (J-) plus an additional 8 seconds if a new test node is created.
4. 4 seconds to add (J-) to entry resulting in (BAJ,J-)
5. 4 seconds to add I to entry resulting in (BAJ,JI-)
6. 4 seconds to add D to entry resulting in (BAJ,JID)

Total learning time here is 24 seconds.

A comparison of learning times that we have just enumerated in the typical unfamiliarized condition (28 seconds) with the condition where the response is familiar (16 seconds) and the condition where the stimulus is familiar (24 seconds) illustrates why EPAM VI produces the pattern of results that has also been found with people as shown in Table 6. As with people, both EPAM III and EPAM VI benefit more from pre-familiarization of responses than they do from pre-familiarization of stimuli. For people the benefit is 1.6 compared to 1.2, for EPAM III it is 1.8 compared to 1.3, and for EPAM VI it is 2.8 compared to 1.9.

Both models of EPAM tend to over-predict the benefits of familiarization. The ratio for people between the completely unfamiliar and completely familiar condition is 1.8, while the ratio for EPAM III is 2.5 and that for EPAM VI is 3.7.⁸ It is possible, however, that for people, even unfamiliar nonsense syllables may have some familiar chunks within them – such as a pair of letters that are familiar or a syllable such as “NUR” that could be seen as the first part of the word “nurse.” Such familiarity would have the same effect upon unfamiliar stimuli as some pre-familiarization, and would speed the ability of people to learn unfamiliar responses.

Table 6

Effects of Stimulus and Response Familiarization

Condition	People	EPAM III	EPAM VI
F-F	1.0	1.0	1.0
U-F	1.2	1.3	1.9
F-U	1.6	1.8	2.8
U-U	1.8	2.5	3.7

Note. Human data are from Chenzoff. (1962) as analyzed in Simon & Feigenbaum (1964). EPAM III data are from Simon & Feigenbaum (1964). Data is the reciprocal of errors in ratio to the F-F condition.

Intralist Similarity of Stimuli and Responses

One surprising regularity is that, although intralist similarity of stimuli effects learning rate, intralist similarity of responses does not (Underwood, 1953). According to EPAM VI, this occurs because only enough about the stimuli has to be learned to differentiate them from each other, but enough has to be learned about the responses to respond with the entire response. Table 7 shows the effects of intralist similarity of the stimuli and of responses. The numbers outside of the parentheses show ratios to the low-low condition while the numbers in parentheses show the actual average number of trials involved.

(Insert Table 7 about here)

EPAM III was run two different ways, normal and “CV-group” in which the initial consonant-vowel were assumed to be an already familiar chunk. An average of the two produced results that are very close to those found in the human data. EPAM VI was also run two different ways, the quick method which did not involve a complete simulation of short term memory and the normal method with a full simulation of short-term memory. The results for the two are slightly different, although the fit with the data and the total number of trials taken is about the same for both versions. Both versions of EPAM VI, like people, had an easier time than

Table 7

Mean Trials to Learn as a Function of Stimulus and Response Similarity

Condition	People	EPAM III		EPAM VI	
		Normal	CV-Group	Quick	With-STM
Low-Low	100 (23.2)	100	100	100 (13.0)	100 (13.4)
Low-Medium	96 (22.4)	88	100	98 (12.8)	97 (13.0)
Low-High	105 (24.4)	91	100	97 (12.6)	97 (13.0)
Medium-Low	110 (25.5)	140	100	118 (15.3)	114 (15.3)
High-Low	132 (30.7)	146	114	124 (16.2)	120 (16.0)

Note. Results expressed as ratio to low-low condition. Where available, the actual number of trials is included in parentheses. Results for people are from Underwood (1953). Results for EPAM III

people with the low-medium list than with the medium-low list even though the two lists consisted of the same exact syllables. (For people the numbers were 96-110; for the two versions of EPAM VI they were 98-118 and 97-114.) Also both versions of EPAM VI, like people, had even bigger spreads when the low-high list is compared with the high-low list. (For people the numbers were 105-132; for the two versions of EPAM VI they were 97-124 and 97-120.)

Both versions of EPAM VI use the same learning mechanisms, the same responding mechanisms, the same discrimination net, and the same declarative memory. Unfortunately, constraints upon our time prevented us from coding all of the simulations using EPAM VI's full version of short-term memory. This simulation, however, shows that when presentation is relatively rapid (in this experiment the stimulus was followed 2 seconds later by the response followed 2 seconds later by the next stimulus) both versions produce very similar results.⁹

Effects of Rehearsal Strategy

EPAM III, like SAL I (e.g. the 2-response vs. 9-response task), was used in a simulation which cast light upon Bower's (1961) one-element model when Gregg and Simon (1967) found that the short-term memory rehearsal strategy called the "one at a time" strategy resulted in one-trial learning although a strategy called the "all-at-once" strategy did not.

These two strategies were tested within a variation of the paired associate paradigm in which study trials alternated with test trials. During a normal paired-associate experiment, subjects must switch their attention to a new stimulus whenever it is presented, because if they do not respond within two seconds their response is counted as incorrect. However during the study stage of a study-test-alternation experiment, subjects do not have to respond to the stimulus at all. Thus, without sacrificing accuracy, they can choose to engage in a "one-at-a-time" strategy where they continue to rehearse a single stimulus-response pair, ignoring other incoming stimuli, until that pair has been learned. Not all subjects engage in this strategy. Others engage in the "all-at-once" strategy in which they turn their attention to new stimulus-response pairs when they are presented during the study stage. Only during the test stage, in which stimuli are presented with no feedback, must all subjects attend to each stimulus when it is presented.

One of the more thorough experiments using this study-test-alternation paradigm was conducted by Gregg, Chenzoff, and Laughery (1963) and reanalyzed by Gregg and Simon (1967) who divided the human subjects by the strategy that they reported, after the experiment, that they had used, either: (1) a "one-at-a-time" strategy in which they tried to learn one or at most two pairs during a trial, or (2) an "all-at-once" strategy in which they tried to learn many or all of the pairs each trial.

The experimental design used by Gregg et al. (1963) varied presentation rate between fast and slow during the study stage and varied whether or not the stimulus-response pairs were replaced with a new pair after a test-stage error. If syllable-pairs were replaced after an error, then only those stimulus response pairs that had been learned in a single trial would be kept for the next trial. Thus, the replacement condition provided a test of whether or not one-trial learning had taken place.

When using the "one-at-a-time" strategy during the study stage, EPAM IV kept a single stimulus-response pair in memory performing repeated calls to its long-term-learning mechanism while rehearsing the pair until, by performing an internal test, it determined that it could correctly respond to the stimulus. On the

Table 8

Results of “one-at-a-time” and “all-at-once” strategies under slow and fast presentation conditions

Condition	People			EPAM VI	
	Overall	One	All	One	All
Slow-Replace	57	100	48	130	16
Fast-Replace	59	68	28	85	9
Slow-Keep	146	—	—	152	152
Fast-Keep	110	—	—	127	127

Note. People’s results are from Gregg et al. (1963) as reanalyzed by strategy by Gregg and Simon (1967). They show the number correct on the 10th trial extrapolated to an n of 20. EPAM VI results are number correct on 6th trail extrapolated to an n of 20.

other hand, when using the “all-at-once” strategy during the study stage. EPAM VI switched to a new stimulus-response pair whenever one was presented, and, as a result, it, like people, was not able to perform the incremental steps necessary to learn a pair in a single trial.

Table 8 shows the results for Gregg et. al’s human subjects and for EPAM VI. The only difference is that the human subjects results are the number correct on the tenth trial, while the EPAM VI results are the number correct on the sixth trial. (The few cases where EPAM VI is able to learn a pair in a single trial using the “all-at-once” strategy result from continuing rehearsal during the 15 second interval in this experiment between the study stage and the test stage.)

Again, EPAM VI, as EPAM III before it, achieves a good fit with the human data. Like the people, EPAM VI is able to learn much better under the replacement condition when it follows a “one-at-a-time”

strategy. The difference for the human subjects between the “one-at-a-time” and “all-at-once” strategy under slow-presentation replacement condition was 100 to 48 for people and 130 to 16 for EPAM VI. The difference under fast presentation conditions for people was 68 to 28 and for EPAM was 85 to 9.

Table 8 also shows that EPAM VI captures the overall difference between slow and fast presentations (both people and EPAM do better with slow presentations than fast) and the overall difference between the “keep” condition and the “replacement” condition. Both EPAM and people do better when responses are not switched after an error.

Incidentally, under the non-replacement condition (marked as “slow-keep” and “fast-keep” in Table 8) the EPAM VI simulated subjects do exactly as well whether they use the “one-at-a-time” or “all-at-once” strategy. Although EPAM VI can learn in a single trial, its learning process is incremental. It can do some steps of learning in one trial (such as forming a chunk for the stimulus) and another step the next (such as creating an entry for uniting the stimulus with the response).

Interactions of Discrimination Net with Declarative Memory

In his classic fan-effect experiment, Anderson (1974) found that entries in declarative memory such as “A hippie is in the park” can be accessed equally from the person (i.e. “hippie”) or from the location (i.e. “park”) even if subjects are only trained to respond from one (i.e. just the person or just the location), not the other.

Anderson (1974) also found that the chunk for a person or a location could point to more than one entry. For example, “hippie” could point to a hippie in the park, another hippie in a church, and another hippie in a school. Similarly a location could point to more than one person. For example there could be a hippie, a lawyer, and a minister in the park.

We have incorporated both of these aspects of declarative memory into EPAM VI. As far as paired-associate experiments are concerned, they mean that the declarative memory entry (B-, JID) can be accessed from both the discrimination net node for the stimulus (“B—”) as well as the discrimination net node for the response (“JID”). Also, a node can point to more than one entry. For example the node for (“G—”) with “Test=2” shown in Figure 6 points to three entries (GXJ, 6), (GX-,7), and (GQJ, 8). SAL III had a similar feature in that a SAL III stimulus node could point to more than one response. These two aspects of declarative memory, borrowed from Anderson’s ACT models (1974, 1983, 1999) allow EPAM VI to simu-

Table 9

Percent backward recall of stimulus as function of intralist similarity

Condition	People	EPAM VI
High-Similarity	72	76
Low-Similarity	34	13

Note. Human data are from Hintzman (1969).

late several experiments whose results depend upon the interaction of the discrimination net with declarative memory.

Backward Recall

None of the previous versions of EPAM or SAL could simulate backward-recall of stimuli from the response. EPAM VI, on the other hand, automatically places pointers to the same entry at both the stimulus node as well as at the response node whenever that entry is studied. The only factor which limits whether a stimulus can be recalled is its completeness in the entry. For example, an entry such as (B-, JID) would only be able to respond that the stimulus started with a “B,” though an entry such as (BAJ, JID) would be able to respond with the entire stimulus.

We tested this feature of EPAM VI by running it on two experiments that Hintzman (1969) conducted with human subjects using two different conditions – low-similarity lists and high similarity lists. The results are shown in Table 9.

For both EPAM and people, where stimuli are more similar to each other they are more easily recalled during a backward-recall trial. The results for people are 72% to 34% and for EPAM VI are 76% to 13%. This result occurs for EPAM VI because similar stimuli must be more fully studied while differentiating them during the learning trials. As a result, high-similarity entries (for examples see the entries of Figure 6) are more

likely to have the entire stimulus in the entry, while low-similarity entries (for examples, see the entries of Figure 5) are likely to only have the first letter of the stimulus in the entry.

The human subjects were better than EPAM VI with backward recall of low-similarity stimuli. Perhaps some of the low-similarity stimuli were already partly familiar to the human subjects, or in other ways pre-familiar. Alternatively, human subjects may do more unmotivated familiarization of stimuli than does EPAM VI.

Hintzman (1969) also measured the distribution of errors by letter position that people made when mispronouncing the stimulus during a backward recall trial and concluded that these errors support the more flexible cue-selection mechanisms that Wynn (1966) incorporated into the WEPAM model, over the mechanisms used by either SAL or EPAM.

Differing Estimates of Recognition and Recall Measures

Experimental paradigms differ regarding the way they determine what the subject has learned. All of the experiments that we have reported thus far have used the “recall” measure in which the subject is asked to recall the entire response given the stimulus (or in the case of backward recall, to recall the entire stimulus given the response). Another measure that can be used is the “recognition” measure in which the subject is presented with either a correct stimulus-response pair (called a “target”) or an incorrect pairing between a stimulus and response (called a “foil”) and the subject is asked, in effect, “Do you ‘recognize’ this stimulus-response pair as being a pair that was studied?”¹⁰

Psychologists have found that subjects score higher with the “recognition” test than with the “recall” test (Postman, Kruesi, & Regan, 1975) and indeed the multiple entries at nodes for SAL III and EPAM VI stimuli allow both models to sometimes recognize even when they don’t recall.

Both EPAM VI and SAL III simulate this difference as shown in Table 10. The key finding is that the probability of recognizing a correct pair is higher than the probability of incorrectly recognizing a foil, even after the response was recalled incorrectly. For example, after a failure to recall the correct response, EPAM VI recognizes the correct stimulus-response pair 66% of the time but only incorrectly recognizes a foil pair 35% of the time. Unlike recall where EPAM VI stops once it finds an entry whose stimulus part matches the stimulus, recognition continues to generate entries until one is found that matches both the stimulus and the response.

Table 10

Recognition Hits and False Alarms with Correct and Incorrect

Recall. Hit indicates that the correct pairing was recognized. False alarm indicates that the foil was (incorrectly) recognized.

Condition	SAL III	EPAM VI
After correct recall		
Probability of a Hit	1.00	1.00
Probability of False Alarm	0.05	0.24
After incorrect recall		
Probability of a Hit	0.41	0.66
Probability of False Alarm	0.22	0.35

Note. SAL III results are from Hintzman (1968).

Second and Third Guesses

People appear to do better if given a second or a third chance to correctly recall a response (Binford and Gettys, 1965; Bower, 1967; Brown 1965). This is also true of EPAM VI and SAL III as shown in Table 11. The numbers above .111 (the chance level for a 9 response list) in the second and third columns show that when given a second or third chance to get it right, the SAL III and EPAM VI simulated subjects often do. The mechanism is simple, if the response on the first entry generated is not correct, the simulated subjects continue to generate entries, and if they run out of entries, they guess from among the possible responses (the numbers 1 through 9).

Time Tags

Inspired by EPAM, Yntema and Trask (1963) theorized that time tags might be associated with each chunk in memory:

The view of human memory that will be adopted here is in many ways similar to that proposed by Feigenbaum (1961). There is however a difference in emphasis....

Table 11

Proportion Correct on First, Second and Third Guesses

Trial number	SAL III			EPAM VI		
	1	2	3	1	2	3
2	.172	.161	.144	0.121	0.121	0.121
3	.350	.256	.287	0.177	0.151	0.130
4	.561	.418	.283	0.262	0.210	0.133
5	.750	.489	.348	0.359	0.275	0.145
6	.822	.531	.467	0.447	0.359	0.146
7	.894	.526	.111	0.540	0.450	0.135

Note. SAL III results are from Hintzman (1968).

A person will be regarded as consisting of two parts, a memory and a processor. For the present purposes the memory may be assumed to contain items of information (chunks, if one prefers) each of which bears a number of tags that describe it and show how it is related to other items in memory.... (p. 65)

Then they designed two experiments to test their theory. Their second experiment used a procedure, later named by Lockhart (1969, p. 42) the “two alternative forced choice (2AFC) recency discrimination task,” in which they presented subjects a series of words on cards which subjects were permitted to study at their own pace. Interspersed between words were test cards which asked subjects to decide which of two words typed on the cards had been presented most recently. Yntema and Trask speculated that subjects might utilize time-tags in order to determine recency. Specifically, they reasoned:


How might the processor be programmed to find the correct response? At first sight the answer seems to be straightforward: the processor should search for items belonging to the given category and choose the most recent. But how is it to decide which is the most recent?

What rules or tests might it use to determine that one item occurred after another. Various hypotheses can be devised, but simplicity is attractive. It is tempting to assume that items in memory can carry time-tags that the processor can examine to determine their relative recencies. (p. 70).

Yntema & Trask indeed found evidence for such time-tags, a finding that was later confirmed by Lockhart (1969) who not only had subjects make relative recency judgements but also absolute recency judgments. We have recharted Lockhart's data putting ratios on the x-axis and percentage correct on the y-axis. The ratios that we have charted were lag from end of list to latter item over lag from end of list to former item. These ratios enabled us to predict Lockhart's data using a simple mathematical model with just one parameter.

Our prediction is based upon the simple idea that people's subjective memory for the time when they

Figure 10. Recency judgments as a function of the recency ratios. The human data shown with X's is from Lockhart (1969). The solid line shows results from EPAM VI..



learned an entry follows a normal curve with a standard deviation of 0.97 when plotted on a log-time scale. Therefore if an event occurred 1 unit of time ago, then the log of the time that the event occurred would be 0. A normal curve can then be plotted having its mean at 0. This normal curve with a mean at zero would have a standard deviation of 0.97 and so about 68% of the values would fall between log-times of -0.97 and 0.97 .¹¹

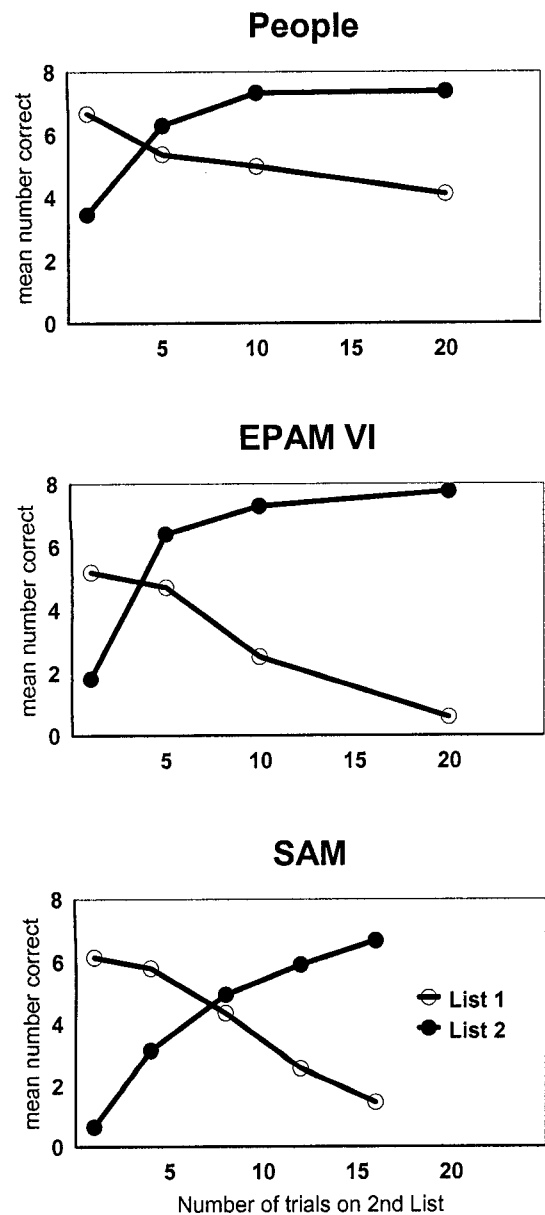
Stored with each EPAM VI entry is the exact time when the entry was last studied. However EPAM is not able to access this time as a precise number. Instead, it accesses an estimate of that time that is drawn randomly from the normal, in log time, distribution around the actual time. In order to make a recency judgment, EPAM VI compares the estimates it has of the time-tags of the two entries.

Figure 10 shows the results of a simulation of EPAM's use of time tags. The "X's" of Figure 10 show the actual data found by Lockhart while the accompanying curve shows the predictions of EPAM VI. These predictions explain 88% of the variance in Lockhart's data.¹² There is only one parameter in this simulation, the standard deviation in log time.

"Unlearning" of Old Responses

EPAM VI's use of time tags enables it to simulate paired-associate experiments where subjects are asked to recall more than one response to a stimulus. For example, Barnes and Underwood (1959) conducted an AB-AC ex-

Figure 11. Responses recalled in the correct column in the AB-AC paradigm as a function of trials on second list. Human data is from Barnes and Underwood (1959). SAM data is from Mensing and Raaijmakers (1988).



periment (i.e. subjects first learned List 1 that associated stimuli A with responses B, and then learned List 2 that associated stimuli A with responses C). Then during an MMFR (“Modified Modified Free Recall”) trial, subjects were asked to write a response to each list next to the stimulus and in the proper column (first column for List 1 responses, second column for List 2 responses).

Figure 11 shows the results for Barnes and Underwood’s subjects, EPAM VI, and the contextual fluctuation version of SAM (Mensink & Raaijmakers, 1988). EPAM VI is able to simulate the gradual drop-off of List 1 responses as List 2 is learned, as shown in Figure 11, although List 1 responses tend to disappear more quickly for EPAM than for people. In order to produce the results in Figure 11, EPAM used its time tags twice – not only to generate items in (imperfect) recency order, but also in order to determine list membership of each item generated. List membership was especially a problem during the early List 2 trials when there were few List 2 responses in the discrimination net. In that situation, EPAM VI was able to use the time tags stored with the entries to determine whether a single response generated belonged in List 1 or List 2.

The SAM model fits the data well because of changing associative strength between the stimulus cue and the interpolated list image without any “unlearning” taking place. EPAM VI also fits the data without having any special unlearning components; it loses access to first list responses when discrimination net learning extends below the nodes where the first list responses reside.

Response Latencies

Just as Feigenbaum was pleasantly surprised when he discovered that EPAM II’s mechanisms explained oscillation and retroactive inhibition, Hintzman (1968) was pleasantly surprised when he discovered that the number of responses in the push-down-stack at a SAL III node correlated with the latency of response when that node was used to generate a response. However, SAL III lacked a mechanism that would explain why this correlation should take place.

EPAM VI, on the other hand, does include a rationale for why the number of entries at a node determines response latencies. When there are multiple entries at a stimulus node, EPAM generates those entries in their order of perceived recency as part of a generate-and-test process which terminates as soon as an entry that matches the stimulus has been found. Generating an entry takes EPAM approximately 250 ms.¹³ Often, several entries must be generated before a matching entry is found.

For example, assume, at the node with “Test=2” in Figure 3a that the time-tags associated with the

entries cause them to be generated in the following order: (XQ-,4), (XF-,3), (X-,1). When “XLG” is next presented and gets sorted to that node, first EPAM VI would generate the entry (XQ-, 4) and would compare its stimulus part “XQ-” with the stimulus that was presented, “XLG.” It would find that the image does not match the actual stimulus and so would generate the next entry (XF-,3) at a time charge of 250 ms and would compare its stimulus part “XF-” with “XLG.” Again there would be no match so EPAM would generate the next entry (X—, 1) at a time charge of an additional 250 ms and would compare its stimulus, “X—” with “XLG.” EPAM would consider this to be a match and so would respond with the response from that entry (i.e. with “1”). Thus the three entries at this node led to an additional 500 msec. of responding latency time when compared with a node in which the first entry generated would be accepted.

Using this mechanism, EPAM VI produces a fairly close approximation of the latency times found by Suppes et al. (1966) in our simulation of that experiment as shown in Figure 12. For both EPAM and people, the response latency times gradually decline when the subject is responding correctly. EPAM VI predicts this result because, during the first trials of a paired-associate simulation, stimuli are often sorted to nodes that contain multiple entries. However, during later trials, the net is further grown so that, in general, stimuli sort to nodes where there is just one entry.

(Insert Figure 12 about here)


EPAM also predicts a finding which Suppes et al. found to be especially difficult for their “three-stage-model” to explain – that the last error for a particular stimulus has a long response latency (1.65 seconds for people, 1.43 seconds for EPAM). According to their three-stage model, the human subjects should be getting closer to the correct response during the second stage, yet latency times did not improve. EPAM VI’s explanation is fairly simple. During the last trial in which a subject responded incorrectly before creating a new branch so that the stimulus would sort to its own node, EPAM was usually responding from a test node to which several entries sorted. such as the “Test=2” node pictured in Figure 3a. But once the last learning occurred, the stimulus would have no further competing entries at its node.

Interlist Similarity and Transfer

Most of the results that we present in this paper were first simulated by SAL III and could not have been simulated by EPAM III. There is, however, one exception in which EPAM VI predicts a finding that was first simulated by EPAM III that could not have been simulated by SAL III.

Figure 12. Mean latency curves conditionalized on errors and successes.

Human data is from Suppes et al. (1966), Session 1.



In this experiment, two of the conditions that were compared were AB-CD and AB-AD. In the AB-CD condition, two completely different lists of paired associates are learned. In the AB-AD condition, the stimuli of the first list were identical to the stimuli of the second. Bruce (1933) found that subjects have an

easier time learning the AB-CD list than the AB-AD list. This pattern of results evidences what is called “negative transfer” since the previous study of the stimuli, proved a negative when the time came to learn the second list.

Negative transfer in EPAM VI depends upon such factors as the time spacing between presentation of the lists and number of lists that have been learned previously. For example, if the lists are presented at widely separated time intervals, there will be little confusion between the time tags of their entries, favoring AB-AD over AB-CD. In order to simulate negative transfer in the Bruce (1933) experiment, we have precisely recreated the time course and previous list learning of that experiment.

Bruce also included a third condition in his experiment labeled AB-CB in which the responses of the first and second lists were identical. EPAM VI, like EPAM III, gets the clear boost during second list learning that pre-familiarization of responses always gives to EPAM.

Table 12 compares the results of EPAM VI with those of Bruce’s (1933) subjects. The columns refer to the number of List 1 trials. For example, in the column labeled “0,” List 1 was not studied at all, just List 2.

Table 12
Number of trials to learn second list as a function of the number of trials of study of the first list.

Condition	People				EPAM VI			
	0	2	6	12	0	2	6	12
AB-AD	100	117	116	109	100	105	110	106
AB-CB	100	115	83	63	100	96	86	80
AB-CD	100	100	108	84	100	100	101	101

Note. Human data is from Bruce (1933). All data is expressed as the ratio to the 0-first-list-study condition.

In the column labeled “12,” the List 1 stimuli were studied for 12 trials. The data shows the ratio of trials to learn this list as compared to the number of trials to learn the list when List 1 was not studied.

The key finding is shown in the “12” column of Table 12. The easiest list to learn, for both EPAM VI and Bruce’s human subjects was the AB-CB list and the hardest was the AB-AD list. Thus for both people and EPAM the high interlist similarity of responses (AB-CB) produced positive transfer when compared to the AB-CD standard condition (63 to 84 for people and 80 to 101 for EPAM VI). While high interlist similarity of stimuli (AB-AD) produced negative transfer when compared to the standard condition (109 to 84 for people and 106 to 101 for EPAM VI).

Fan Effects

One of the few areas related to the paired-associate paradigm that has progressed over the past 25 years has been the “fan-effect” recognition-paradigm which was developed, explained, and named by John R. Anderson. We are able to explain these effects partly because we have adopted, as noted at the beginning of this paper, a model of declarative memory derived from Anderson’s ACT models.

There are, however a few minor differences between our explanation using EPAM VI and his explanation using ACT-R:

1. We explain the effect through serial processes each having a time requirement while he explains the effect through parallel processes which assume that activation level equates with time.
2. To a certain extent our time-tags play a similar role in our model that activation plays in his model. Just as activation results in quicker responding in ACT-R, more recent time tags result in quicker responding in EPAM. This occurs in EPAM because multiple entries at a node are generated in approximately the perceived recency order of their time tags. On the other hand, recency is only one aspect that determines activation in ACT-R. It may be that more of the determinants of activation could be imported into EPAM VI from ACT-R. (We look forward to a future combination of EPAM with ACT-R that could explain those phenomena previously explained by both.)

In Anderson’s (1974) experiment which first discovered the fan-effect, he presented people with various sentences for study such as “The hippie is in the park,” “The lawyer is in the bank,” and “The lawyer is in the park.” Then he trained subjects to answer questions such as, “Where are the hippies?” or “Who is in the bank?” Finally he tested subjects using the standard technique of the recognition paradigm, which in effect had

them recognize whether or not they had studied target sentences such as “The lawyer is in the park” or foil sentences such as “The hippie is in the bank.” He told them to answer as quickly as they could and he measured their latency of response.

His main finding was that the more places a person was, or the more people were in a location, the longer it took people to recognize the sentences. If a person was in three places, then that person was said to have a fan of 3. If a location had two different people, then that location was said to have a fan of 2.

In order to simulate this effect with EPAM VI, we have fitted one parameter to each experiment: the fixed time to recognize and respond. We have set the fixed time to respond at 1.11 sec. in our simulation of Anderson’s (1974) experiment and at 1.26 seconds in our simulation of Anderson and Reder’s (1999) experiment. It may be that slightly different experimental apparatus or responding procedures causes variations in overall latency times. To this fixed time gets added the occasional extra 250 ms if the system generates and tests two entries while responding and the extra 500 ms if the system generates and tests three entries. We suppose that differences in experimental equipment or differences in subjects’ abilities could result in overall differences in reaction times. For all of the other experiments reported in this paper, we have used the exact same value for each of EPAM’s parameters.

Essentially EPAM VI performs the task in the following fashion:

1. EPAM is presented with a sentence.
2. It flips a coin and based upon the coin toss chooses to generate entries at the person or at the location.
3. It then proceeds to generate entries testing each entry to see if it matches the person and location. Each entry after the first takes 250 msec. to generate and test.
4. If it has generated all of the entries, but has not found a match, it responds that it does not recognize the sentence.
5. During the test stage, subjects are given feedback. EPAM can engage in learning after a target is presented both after a correct response and also after an error. These learning procedures are generally the same that take place in a paired-associate experiment.
6. After a correct response when EPAM’s long-term-learning is not busy, EPAM’s over-learning parameter “*SAL-C*” (which is equal to 30%) determines whether EPAM will study a target.

Thus 15% of the time EPAM will further discriminate the person in the net and 15% of the time it will further discriminate the location. If it discriminates the person, that person is moved to a different node where the fan, the next time that target is presented, will just be 1 and where other targets having the same person will, at first, not be recognized.

7. After an error, when EPAM’s long-term-learning is not busy, EPAM will study the target by updating its time-tag and restoring its pointers to the entry from the person and location nodes.

Step 6 results in overlearning already familiar words, such as “hippie.” Although EPAM is still mainly a model of nonsense syllable, not conceptual learning, we suspect that people are somehow learning in a visual-conceptual realm by adding features to their mental picture of the hippie that is in the bank. Perhaps they are envisioning the hippie as male, adding a goatee or putting beads about his neck. Then the next time they hear the word “hippie” they conjure up the image of this hippie and find the entry (Hippie, Bank).

Table 13 and Table 14 compare the response latencies of human subjects (from Anderson 1974) with those of ACT-R (Anderson & Reder, 1999) and EPAM VI. Both EPAM VI and ACT-R provide close fits to the human data.

Table 13

Mean Reaction Times (in Seconds) for TARGETS for Fans of 1, 2, and 3

Location	People			ACT-R			EPAM VI		
	Fan			Fan			Fan		
	Fan per Person	Fan per Person	Fan per Person	Fan per Person	Fan per Person	Fan per Person	Fan per Person	Fan per Person	Fan per Person
	Fan-1	Fan-2	Fan-3	Fan-1	Fan-2	Fan-3	Fan-1	Fan-2	Fan-3
1	1.11	1.17	1.22	1.11	1.15	1.18	1.12	1.16	1.21
2	1.17	1.20	1.22	1.15	1.22	1.26	1.17	1.20	1.24
3	1.15	1.23	1.36	1.18	1.26	1.33	1.23	1.25	1.30

Note. Human data are from Anderson (1974). Act-R data are from Anderson and Reder (1999).

Table 14

Mean Reaction Times (in Seconds) for FOILS for Fans of 1, 2, and 3

Location Fan	People			ACT-R			EPAM6		
	Fan per Person			Fan per Person			Fan per Person		
	Fan-1	Fan-2	Fan-3	Fan-1	Fan-2	Fan-3	Fan-1	Fan-2	Fan-3
1	1.20	1.22	1.26	1.22	1.27	1.31	1.12	1.22	1.32
2	1.25	1.36	1.29	1.27	1.32	1.36	1.22	1.31	1.40
3	1.26	1.47	1.47	1.31	1.36	1.39	1.37	1.41	1.51

Note. Human data are from Anderson (1974). Act-R data are from Anderson and Reder (1999).

Both show the fan-effect in Figure 15 that the more types of person in a location, or the more locations associated with a type of person the longer the recognition time. For example, an examination of the diagonals shows that when the fans go up from 1 to 2 to 3 at both the person and the location, the recognition latencies go up from 1.11 to 1.20 to 1.36. For ACT-R they go from 1.11 to 1.22 to 1.33 and for EPAM VI they go from 1.12 to 1.20 to 1.30.

Both show the overall effect for foils shown in Table 14 that foils take longer to reject than targets take to be recognized in that the numbers are generally higher in this table than they are in Table 13. An examination of the diagonals in Table 14 shows that when the fans go up from 1 to 2 to 3, the latencies go up from 1.20 to 1.36 to 1.47 for people, from 1.22 to 1.32 to 1.39 for ACT-R, and from 1.12 to 1.31. to 1.51 for EPAM VI showing that the time to reject a foil increases as a function of fan.

Anderson and Reder (1999) conducted a more complex fan-effect experiment in which they varied the rate of presentation of stimuli and observed the effect that this had upon response times. Target stimuli were divided into a facilitation class which were presented more frequently than the other stimuli, an interference class whose sentences shared either a particular person or a particular location with the facilitation class,

Table 15

Mean Latencies in seconds & Error Rates in percent for 2 and 4 Fans

Condition	People		Act-R		EPAM6	
	Fan-2	Fan-4	Fan-2	Fan-4	Fan-2	Fan-4
Targets						
Facilitation						
Mean latency	1.33	1.42	1.36	1.41	1.33	1.40
Error rate	6.2	6.0			0.3	1.1
Interference						
Mean latency	1.57	1.58	1.53	1.62	1.36	1.47
Error rate	10.5	14.0			2.5	3.3
Suppression						
Mean latency	1.43	1.51	1.43	1.49	1.38	1.48
Error rate	4.9	7.8			0.9	1.6
Control						
Mean latency	1.42	1.53	1.43	1.49	1.38	1.47
Error rate	5.5	8.9			1.0	1.8
Foils						
High						
Mean latency	1.51	1.60	1.55	1.60	1.43	1.56
Error rate	7.5	9.9			0.1	0.3
Mixed						
Mean latency	1.58	1.67	1.59	1.64	1.46	1.60
Error rate	8.3	13.8			0.1	0.2
Low						
Mean latency	1.59	1.70	1.62	1.68	1.47	1.65
Error rate	11.7	14.4			0.1	0.2

Note. Human data and ACT-R data are from Anderson and Reder (1999).

a suppression class whose sentences shared either a particular person or a particular location with the interference class but not with the facilitation class, and a control class which did not share person or location with the other classes. Foils as well were divided into classes depending upon the frequency with which their person or location was presented in the targets.

EPAM captures the overall pattern of results shown in Table 15. For example, that facilitation targets

and high foils are responded to most quickly and that, among targets, interference targets have the highest error rate.

EPAM's over-learning process (step 6, above) is largely responsible for this pattern of results in this experiment. Since facilitation targets are presented more frequently than the other sentences, they are more likely to be overlearned, putting them, for a while, as the only entries at their own nodes, thus making responding to them and to foils that are highly related, more rapid, and making responding to the interference targets less accurate.

EPAM does not produce as much variation or as many errors as are found with human subjects. For example, although it produces more interference errors than suppression or control errors, it does not produce slower interference latencies. Perhaps there are features of ACT-R, not yet incorporated into EPAM VI, that could improve its fit with these data.

Probability Matching

Time tags, however, do not always improve EPAM VI's fit with the human data. In a paired-associate simulation modeled loosely upon the procedures of Voss et al. (1959) in which two different responses were associated with each stimulus and were reinforced at various rates of reinforcement. Voss et al.'s subjects overmatched high input probabilities, SAL directly mirrors the input probabilities, and EPAM VI normally undermatches high input probabilities as shown in the "normal" column of Table 16.

These simulations measure the proportion of responses by SAL I and EPAM VI after the discrimination nets have become elaborated sufficiently to discriminate each stimulus from all of the others. If SAL responds correctly, it keeps the same response, if not, there is a fixed probability defined by its learning parameters which causes it to replace the response with the just reinforced alternative.

At this point in the experiment, EPAM VI usually has an entry for each response at the appropriate discrimination net node. EPAM chooses its response based upon the time tags of these entries. Whenever EPAM responds incorrectly, it updates the time-tag of the correct entry. If these time tags were perfectly accurate measures, then EPAM would engage in perfect probability matching, like SAL. However, the time tags tend to become less reliable with the passage of time allowing the less reinforced response to spontaneously pop up even when it was not the last one refreshed.

This undermatching of high input probabilities should not be seen as a refutation of EPAM VI's time

Table 16

Input Probability and Response Proportion

Input Probability	People	SAL I	EPAM VI	
			Normal	Revised
0.1	0.08	0.10	0.19	0.07
0.2	0.06	0.21	0.29	0.17
0.3	0.15	0.30	0.37	0.28
0.4	0.25	0.38	0.44	0.38
0.5	0.46	0.50	0.50	0.50
0.6	0.67	0.62	0.56	0.62
0.7	0.78	0.70	0.63	0.72
0.8	0.83	0.79	0.71	0.83
0.9	0.90	0.90	0.81	0.93
1.0	1.00	1.00	1.00	1.00

Note. SAL I results are from Hintzman (1968). Results for people are from estimated from the Figure 1 data for the 12th block of 10 trials in Viss, Thompson, and Keegan (1957).

tags. Voss et al.'s subjects were given unusual paired-associate instructions: "Your task is to say the response word which you think will be presented with the stimulus word on all subsequent presentations of the same stimulus word" (p. 390). Perhaps their strategy deviated from the normal paired-associate strategy.

If EPAM VI were programmed with a revised strategy that caused it to always study after a successful response (*SAL-C* set at 100%) and to only study after error 25% of the time. Then EPAM would also tend to overmatch these high input probabilities as shown in the "Revised" column of Table 16.

Exponential Decay

Exponential decay occurs with radioactive elements. The half life is the time period during which half

of the radioactive particles will spontaneously leave the element. Following Hintzman, we have used an exponential function to define the probability of a pointer to an entry disappearing from a discrimination net node. Specifically, we have theorized, borrowing from SAL III (Hintzman, 1968), that whenever there is more than one pointer to an entry of a certain type at an EPAM VI node, only the first pointer that was added to the node is stable. All of the others tend to spontaneously disappear with a half life of about 38.5 hours.

Since only pointers to the newer entries, not the pointer to the oldest one, tend to decay, this theory allows us to have exponential decay without violating Jost's law (Simon, 1966) that "If two associations are of equal strength but of different age, the older diminishes less with time" (McGeoch, 1942, p. 140).

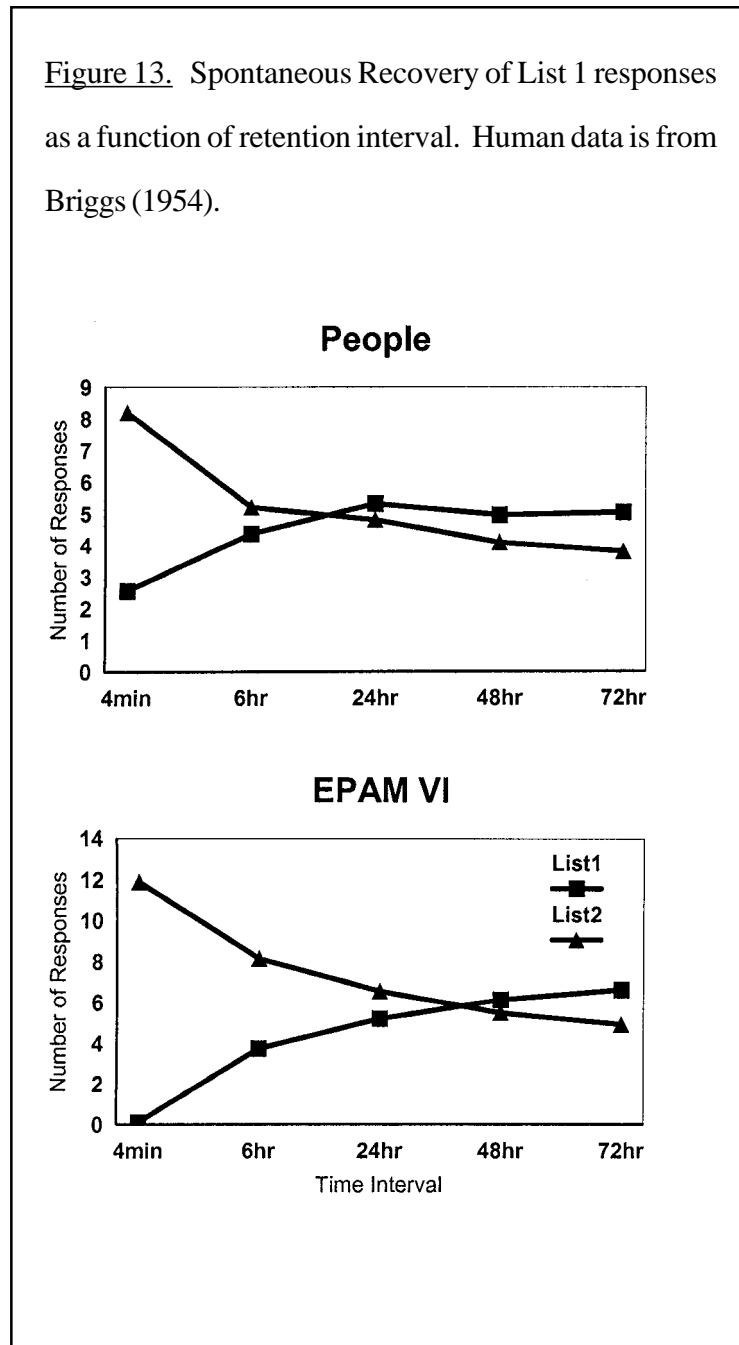
This decay, incidentally, allows the system to respond more rapidly to stimuli and may actually be an efficient way of preserving rapid responding to old items. In order to make a pointer to a new entry permanent, EPAM must move it to a new node, or, if that is impossible, EPAM must engage in sufficient overlearning so that the new pointer replaces all older pointers at the stimulus node.

In one of the few paired-associate studies to measure forgetting over a several day period, Briggs (1954) taught his subjects the AB list to a criteria of one perfect trial and then 24 hours later taught them the AC list, again to a criteria of one perfect trial. Then at various intervals after teaching the AC list, ranging from 4 minutes to 72 hours, he presented subjects with the stimuli in order to find if they would respond with the first list response or the second list response. He found that subjects generally respond with the second list response after four minutes, but that as the time period extends, the proportion of second list responses tends to fall while the proportion of first list responses tends to recover until, by 72 hours, subjects respond more frequently with first list, than with second list responses.

There are two factors that cause the recovery of first list responses in EPAM VI:

1. The imprecise estimate of time tags causes the entries to be confused so that the first list entries might actually be generated before the second list entries. This factor alone, however, could not cause the first list entries to eventually become more common than the second list entries.
2. Where the entries of second list responses co-exist at nodes with first list entries, those entries will tend to be unstable and to disappear from the nodes at an exponential rate with a half-life of 38.5 hours. This permits the first list entries to actually become more common than the second list entries.

Figure 13. Spontaneous Recovery of List 1 responses as a function of retention interval. Human data is from Briggs (1954).



Briggs conducted his experiment using common adjectives, not nonsense syllables. He assumed that subjects already had learned an old response to each stimulus adjective and determined what that might be by asking subjects to give the first word that came to mind upon seeing each of the stimulus adjectives. In order to give EPAM such old responses, we ran the experiment in AB-AC-AD fashion. First we ran one trial of AB learning (designated as List 0) so that the subjects would learn some of these old responses to the stimuli. Then, a simulated one year later, EPAM learned AC (List 1) to a criteria of one consecutive perfect trial. Then

a simulated one day later, EPAM learned AD (List 2) to a criteria of one consecutive perfect trial. Following that, trials occurred at intervals that varied from 4 minutes to 72 hours in which EPAM responded with the first response generated at the stimulus node (Briggs' subjects were asked to respond with the "first word that comes to mind" and if possible with one of the responses that was assigned to each stimulus word).

EPAM produces a good fit with the human data, shown in Figure 13. For both people and EPAM VI, the List 2 responses gradually fall off while the List 1 responses spontaneously recover. For both EPAM VI and people, the List 1 responses have passed the List 2 responses by the end of the 72 hour interval.

EPAM VI's data would be much closer if it were to make more List 1 substitutions after 4 minutes of learning. Perhaps this difference is due to the use of words by Briggs as his stimuli and responses while EPAM used nonsense syllables. Alternatively, it may indicate that there is another factor involved that is not captured by the EPAM model. One possibility could be that every pointer position at a discrimination net node has a different exponential forgetting rate. A third pointer could be forgotten with a half life in minutes, not hours, a fourth with a half life in seconds. We could fit this data much more closely if we were to introduce such additional parameters.

Slot Filling

As Simon (1972) noted, EPAM III would not be able to explain a class of paired-associate experiments, reviewed by Bower (1972), involving visualizable relations between meaningful stimuli and responses. EPAM VI, on the other hand, can explain these results by using the theory of expertise developed when EPAM IV (Richman et al., 1995) simulated the ability of a lab-trained mnemonist to correctly repeat up to 100 digits presented one per second by adding them, and redundant information about them, to slots of a retrieval structure built, internally, for that purpose.

In this paper we hold that people are "experts" when it comes to using language and that they can fill syntactic slots with meaningful information just as quickly as our lab-trained mnemonist could fill his retrieval structure slots with digits.

In these experiments subjects were asked to pair a concrete noun stimulus (such as "whale") with a concrete noun response (such as "cigar"). The fastest learning times occurred when subjects either visualized the stimulus and response in an interactive scene (such as the cigar in the whale's mouth) or were given a meaningful sentence ("Whale smoking cigar."). The slowest learning times occurred when subjects were given

a nonsense sentence (“Whale trotting cigar.”).

We suppose that in the interactive scene condition the subjects were able to use a slotted-schema for a whale which had a slot for mouth in which certain objects, such as cigars could be placed. In the meaningful sentence condition we assume that the verb “smoking” had syntactic slots in which the whale and the cigar could be quickly placed. However, in the nonsense sentence condition, subjects used the same rote-learning strategy without slots that they would have used when pairing nonsense syllables.

Earlier we discussed the basic EPAM time-constraints involved when two familiar nonsense syllables (or familiar words) are associated in rote-learning fashion. To summarize, that learning process involves the following steps, requiring 14 seconds:

1. 4 seconds to form entry ()
2. 4 seconds to add (WHALE) to entry resulting in (WHALE,)
3. 1 second to create a link from the chunk for WHALE to the new entry.
4. 4 seconds to add (CIGAR) to entry resulting in (WHALE,CIGAR)
5. 1 second to create a link from the chunk for CIGAR to the new entry.

However, if EPAM VI is given a syntactic structure such as (SMOKING,_,_) then we assume that the learning process takes just 6.5 seconds:

1. 4 seconds to form a new entry as an instantiation of the slotted-schema (SMOKING,_,_)
2. ¼ seconds to add (WHALE) to the subject slot resulting in (SMOKING,WHALE,_)
3. 1 second to create a link from the chunk for WHALE to the new entry.
4. ¼ seconds to add (CIGAR) to entry resulting in (SMOKING,WHALE,CIGAR)
5. 1 second to create a link from the chunk for CIGAR to the new entry.

Thus the total time to learn to associate whale with cigar using a syntactic structure is just 6.5 seconds. However, there is an additional factor that enters into this equation which suggests that slots once filled, may not always stay filled. We discovered (Richman et al., 1995) that our lab-trained mnemonist appeared to lose about 21% of the information that he placed in the retrieval structure slots, but was able to usually recover the missing information because he remembered information redundantly. Not yet having corroborating information that similar amounts of information get lost from syntactic slots, we have guessed that this forgetting rate is a universal characteristic of slots.

We have run EPAM VI in a paired-associate task using, approximately, the procedure described by Bower. EPAM like the human subjects had one 5-second exposure to each of 20 pairs. For both EPAM and the human subjects, the measured result was the immediate recall after 1 learning trial. While the human subjects learned to associate concrete nouns, EPAM learned to associate prefamiliarized nonsense syllables.

For EPAM, the only difference between the two conditions was that in the “Imagery” condition EPAM joined the stimulus and response within an entry having two slots while in the “Control” condition, EPAM joined the stimulus and response within an entry having no slots. The results for EPAM and the human subjects are displayed in Table 17.¹⁴ These results are quite close indicating that we are in the right ballpark in our estimates of the time that is saved by using slotted schema.

Slots may even be used occasionally during rote learning experiments when nonsense syllables are paired. For example, the nonsense syllables that Bugelski (1962) used in his experiment suggested real words or phrases to some of his subjects. For example, as Bugelski (1962, p. 411) noted, “CEZ-MUN becomes SAYS MAN or C’EST MAN or SEND MONEY.” It may be that subjects of paired-associate simulations sometimes use slotted structures to join stimulus with response.

Table 17

Percentage of immediate recall after 1 trial of learning with a 5 second presentation rate.

Condition	People	EPAM6
Imagery	0.68	0.63
Control	0.37	0.35

Note. Human data is estimated from list 1 of Figure 3.1 of Bower (1972).

Discussion

We have shown that EPAM VI captures 20 regularities from the paired-associate paradigm. Simply by themselves, discrimination nets allow EPAM VI to explain:

1. Intralist similarity. EPAM VI, like all its predecessors, explains Hintzman's (1969) finding that more errors are made when the stimuli are similar to each other.
2. Subject protocols. EPAM VI, like SAL I and EPAM II can explain the stimulus-generalization, perseveration, and oscillation that is apparent in people's protocols. Unlike its predecessors, EPAM VI's time tags also enable it to explain the tendency for people to return to making some errors even after the perfect trial. Unlike WEPAM (Wynn, 1966), however, EPAM VI can not explain confusions between pairs of trigrams which have the same letters in different positions.
3. Number of response alternatives. EPAM VI, like SAL I, explains the differences found by Hintzman (1967) between human subjects' pre-criterion curves for 2-response lists and 9-response lists.
4. Interlist stimulus similarity and retroactive inhibition. EPAM VI, like SAL II and EPAM II, explains the effect reported by McGeoch and Irion (1952) that the more similar first list stimuli are to second list stimuli, the more first list responses will be forgotten when the second list is learned.

With the addition of SAL II's overlearning component, EPAM VI can explain an additional aspect of retroactive inhibition.

5. Degree of interpolated list learning and retroactive inhibition. EPAM VI, like SAL II, explains the complex interactions found by Thune and Underwood (1943) of degree of interpolated list learning upon total errors and intrusion errors made when relearning a first list.

With the addition of the chunking mechanisms and rehearsal strategies developed in EPAM III, EPAM VI can explain:

6. Time required to learn. EPAM VI explains Bugelski's (1962) finding that total learning time is constant no matter the rate of presentation of a list of nonsense syllables.
7. Prefamiliarization of stimuli and responses. EPAM VI, like EPAM III, explains Chenzoff's (1962) finding that advance familiarity with stimuli is not nearly as helpful as advance familiarity with responses.
8. Intralist similarity of stimuli and responses. EPAM VI, like EPAM III, explains Underwood's (1953) finding that intralist similarity of stimuli slows learning, but intralist similarity of responses

does not.

9. Effects of strategy in replacement and non-replacement conditions. EPAM VI, like EPAM III, explains why a learning strategy difference in an experiment by Gregg et al. (1963) affects the results of a replacement experiment in which responses are changed whenever the subject responds incorrectly.

With the addition of the declarative memory of the ACT models (Anderson, 1974, 1983), EPAM can explain:

10. Backward recall. Unlike its predecessors, EPAM VI explains Hintzman's (1969) finding that the higher the intralist similarity of the stimuli, the more often a stimulus can be correctly recalled, given the response.
11. Recognition can succeed even though recall fails. EPAM VI, like SAL III, is able to capture the finding that a stimulus-response pair can often be recognized even though the response, given just the stimulus, cannot be recalled. This is similar to results found by Postman, Kruesi, and Regan (1975).
12. People do better than chance on second and third guesses. EPAM VI, like SAL III, explains the finding that people do better than chance with second or third guesses (Binford and Gettys, 1965; Bower 1967; Bergman 1966; Brown 1965).

With the addition of time tags first theorized by Yntema and Trask (1963) and estimated here to have a standard deviation of 0.97 when plotted on a log-time scale, EPAM VI can explain.

13. Recency judgments in a forced choice task. Our one-parameter model of time tags explains 88% of the variance found in Lockhart's (1969) data.
14. Unlearning of old responses. EPAM VI, like SAL III, tracks the gradual unlearning of first list responses found by Barnes and Underwood (1959) when a second list is learned.
15. Response latencies over the course of a paired associate experiment. EPAM VI, like SAL III, explains the gradual reduction in response latencies over the course of a paired associate experiment found by Suppes et al. (1966). EPAM VI also explains a finding, troubling for Suppes et al.'s three-stage mathematical model, that the last error for a particular stimulus has a long response latency.
16. Interlist similarity and transfer. EPAM VI, like EPAM III but unlike SAL III, explains the positive transfer found by Bruce (1933) when responses of two successive lists are the same and the

negative transfer when the stimuli of two successive lists are the same.

17. Fan-effects. EPAM VI explains the fan-effects found by Anderson (1974) and Anderson and Reder (1999).

18. Probability matching. EPAM VI, like SAL I, explains the probability matching behavior found by Voss et al. (1959) that when two responses are associated with the same stimulus, then subjects' responses will roughly match the reinforcement proportions of those responses. As a result of its time tags, however, EPAM VI normally tends to undermatch probabilities at high input probability levels while Voss et al.'s (1959) subjects tend to overmatch. A different strategy programmed into EPAM VI, motivated by Voss et al.'s unusual paired associate instructions, gets EPAM VI to overmatch high input probabilities.

With the addition of exponential decay of newer pointers to entries at a discrimination net node, EPAM can explain:

19. Spontaneous recovery of old responses. EPAM VI, like SAL III, simulates the spontaneous recovery of old responses found by Briggs (1954).

With the addition of the theory of slotted structures developed by Richman et al. (1995) EPAM VI explains:

20. Visualizable relations between stimulus and response. EPAM VI, unlike EPAM III, explains why it takes subjects less time to learn to associate "whale" with "cigar" if they are given the meaningful sentence "whale smoking cigar" rather than the nonsense sentence "whale trotting cigar" (Bower, 1972; Simon, 1972).

A version of the SAM model with contextual fluctuation processes (Mensink & Raaijmakers, 1988) also explains two of these experiments, perhaps because it shares many features with EPAM: (1) both EPAM VI and SAM are closely related to Atkinson & Shiffrin's (1968) two-stage model with transfer to long-term memory occurring during short-term memory rehearsal; (2) both models have entries in long term memory which hold information about the associated stimulus and response; (3) both find the stimulus-response entry through a matching process although EPAM VI's process just matches the stimulus locally to those entries sorted to by the discrimination net while SAM's process matches the stimulus globally to all entries in memory; (4) both models are imperfectly able to distinguish recency, EPAM VI through time tags while SAM through contextual fluctuations; (5) while SAM does not build discrimination nets, it does gradually add features to the stimulus representation at the entry so that the stimuli are more easily discriminated from each other. It may be

that SAM is just a global-matching version of EPAM VI *or* EPAM VI is simply a discrimination-net version of SAM. It remains to be seen, however, whether SAM could simulate the wide range of paired associate data simulated by EPAM VI.

Bower's (1961) one-element model makes excellent mathematical predictions when it does apply and with just one parameter. But it does not include a consistent theory of when it will apply and when it will not. EPAM VI, on the other hand, supplies that theory. EPAM VI predicts that the one-element model will not apply when stimuli are similar unless there are only two possible responses (Hintzman, 1967). It also predicts that the one-element model will only apply in paired-associate replacement experiments when human subjects use a "one-at-a-time" strategy (Gregg & Simon, 1967).

The fits of EPAM VI with the human data in these many simulations are not perfect; there is clearly room for improvement. It would be helpful, for example, for future revisions to include WEPAM's (Wynn, 1966) idea that tests could look for the presence of letters, not just for the contents of a list position. Then EPAM would better capture the transposition errors that appear in human protocols and the errors by serial position when people do the backwards-recall task (Hintzman, 1969).

It would also prove helpful to work out how people do the AB-AC-AB task using the precise conditions of the Thune and Underwood (1943) experiment. Perhaps an improved model of context, as Hintzman (1968) suggested, would help.

Slotted-schema are relatively new additions to EPAM and their effects have not been worked out fully. It would prove helpful if a future version of EPAM were to include the automatic slot creation procedures that have worked well for EPAM VI's close cousin CHREST (Gobet, 1993, 1996; Gobet & Simon, 1998, 2000) when simulating chess memory and perception. It may be that people often search for slotted schema in order to give meaning (as Bugelski's [1962] subjects appeared to do) to the stimulus-response association.

EPAM VI combines the two most successful models of this paradigm thus far, SAL III and EPAM III. It then sprinkles in the declarative memory worked out by Anderson (1974), the slotted schema worked out by Richman et al. (1995), and the recency time tags postulated by Yntema and Trask (1963) and estimated, here (using Lockhart's [1969] data) to have a standard deviation from accuracy of 0.97 in log time.

EPAM VI continues to simulate the serial anticipation data earlier simulated by EPAMI (Feigenbaum

& Simon, 1962), the context effects in letter perception data earlier simulated by EPAM IIIA (Richman & Simon, 1989), the expert memory for digits data earlier simulated by EPAM IV (Richman et al., 1995), and the classification learning under strategic control data earlier simulated by EPAM V (Gobet et al., 1997). EPAM VI thus provides an internally consistent explanation of a wide variety of phenomena, and is not just a model of the paired associate paradigm. Current reports are being submitted for publication about its simulations of a wide variety of experimental paradigms including classification learning, distractor task, and articulatory loop paradigms.

EPAM VI's discrimination net cousins CHREST (Gobet, 1993, 1996; Gobet & Simon, 1998, 2000), CLS (Hunt, Marin, & Stone, 1966), FREPAM (Allen, 1965), and CYRUS (Kolodner, 1986) have provided some of the best existing simulations of chess perception and memory, non-incremental classification learning, free recall of nonsense syllables, and reconstructive episodic memory, respectively.

EPAM VI shows that it is indeed possible to distill a theory from the half-century of paired-associate research, a theory that not only explains research of the past, but that also can cast light on the problems that we investigate today. In the *1979 Annual Review of Psychology* (Simon, 1979, p. 378) one of us said, "The safest conclusion at the present time is that human LTM can be represented as a node-link memory with an EPAM-like index, but that various alternatives are still open for the detailed structure and organization of that memory." Later Feigenbaum and Simon (1984) repeated that claim. And today we repeat it here.

References

- Allen, M.M. (1965). The effect of formal similarity in three verbal learning tasks. Doctoral Dissertation. University of California, Berkeley.
- Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. Cognitive Psychology, *6*, 451-474.
- Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. Journal of Experimental Psychology: General, *128*, 186-197.
- 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.
- Atkinson, R.C. & Shiffrin, R. (1968). Human memory: A proposed system and its control processes. In The psychology of learning and motivation: Advances in research and theory, Vol. 2 (ed. K.W. Spence and J.R. Spence), pp. 89-195. New York, NY: Academic Press.
- 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.
- Barnes, J.M., & Underwood, B.J. (1959). "Fate" of first-list associations in transfer theory. Journal of Experimental Psychology, *58*, 97-105.
- Bilodeau, E.A., & Schlosberg, H. (1951). Similarity in stimulating conditions as a variable in retroactive inhibition. Journal of Experimental Psychology, *41*, 199-204.
- Binford, J.R., & Gettys, C. (1965). Nonstationarity in paired-associate learning as indicated by a second guess procedure. Journal of Mathematical Psychology, *2*, 190-195.
- Bower, G. H. (1961). Application of a model to paired-associate learning. Psychometrika, *26*, 255-280.
- Bower, G.H. (1967). Notes on a descriptive theory of memory. In D.A. Kimble (Ed.), Learning, remembering, and forgetting: Proceedings of the second conference. New York: New York Academy of Science..
- Bower, G.H. (1972). Mental imagery and associative learning. In L.W. Gregg (Ed.), Cognition in learning and memory. New York: John Wiley.

Briggs, G. E. (1954). Acquisition, extinction, and recovery functions in retroactive inhibition. Journal of Experimental Psychology, *47*, 285-293.

Brown, J. (1965). A comparison of recognition and recall by a multiple-response method. Journal of Verbal Learning and Verbal Behavior, *4*, 401-408.

Bruce, R.W. (1933). Conditions of transfer in training. Journal of Experimental Psychology, *16*, 343-361.

Bugelski, B.R. (1962). Presentation time, total time, and mediation in paired-associate learning. Journal of Experimental Psychology, *63*, 409-89.

Bush, R.R., & Mosteller, F.A. (1959). A comparison of eight models. In R.R. Bush & W.K. Estes (Eds.), Studies in mathematical learning theory. Stanford: Stanford University Press.

Calfee, R.C., & Atkinson, R.C. Paired-associate models and the effects of list length. Journal of Mathematical Psychology, *2*, 254-266.

Chappell, M. and Humphreys, M.S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. Psychological Review, *101*, 103-128.

Charness, N. (1976). Memory for chess positions: Resistance to interference. Journal of Experimental Psychology: Human Learning and Memory, *2*, 641-653.

Chase, W.G., & Simon, H.A. (1973). Skill in Chess. American Scientist, *61*, 393-403.

Chenzoff, A.P. (1962). The interaction of meaningfulness with S and R familiarization in paired associate learning. Unpublished doctoral dissertation, Carnegie Institute of Technology.

Feigenbaum, E. A. (1959). An information processing theory of verbal learning (Report No. P-I 857). Santa Monica, CA: RAND Corporation.

Feigenbaum, E.A. (1961). The simulation of verbal learning behavior. Proceedings of the Western Joint Computer Conference, *19*, 121-132.

Feigenbaum, E.A. (1963). The simulation of verbal learning behavior. In E.A. Feigenbaum & J. Feldman (Eds.), Computers and thought. New York: McGraw-Hill.

Feigenbaum, E.A., & Simon, H.A. (1962). A theory of the serial position effect. British Journal of Psychology, *53*, 307- 320.

Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. Cognitive

Greeno, J.G. (1964). Paired-associate learning with massed and distributed repetitions of items. Journal of Experimental Psychology, *67*, 286-295.

Gregg, L.W., Chenzoff, A.P., and Laughery, K.R. (1963). The effect of rate of presentation, substitution, and mode of response in paired-associate learning. American Journal of Psychology, *76*, 110-115

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. (1967). Some tests of a discrimination net theory: Paired associate learning as a function of stimulus similarity and number of responses. Journal of Verbal Learning and Verbal Behavior, *6*, 809-816.

Hintzman, D.L. (1968). Explorations with a discrimination net model for paired-associate learning. Journal of Mathematical Psychology, *5*, 123-162.

Hintzman, D.L. (1969). Backward recall as a function of stimulus similarity. Journal of Verbal Learning and Verbal Behavior, *8*, 384-387.

Science, *8*, 305-336.

Gobet, F. (1993). A computer model of chess memory. Proceedings of the 15th Annual Meeting of the Cognitive Science Society, 463-468.

Gobet, F. (1996). Discrimination nets, production systems and semantic networks: Elements of a unified framework. Proceedings of the second international conference on the learning sciences. Evanston IL: Northwestern University.

Gobet, F., Lane, P.C.R., Croker, S., Cheng, P.C., Jones, G., Oliver, I., & Pine, J. (2001). Chunking mechanisms in human learning. Trends in Cognitive Science, *5*, 236-243.

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.

Gobet, F., & Simon, H.A. (1998). Expert chess memory: Revisiting the chunking hypothesis. Memory, *6*, 225-255.

Gobet, F., & Simon, H.A. (2000). Five seconds or sixty? Presentation time in expert memory. Cognitive Science, *24*, 651-682.

- Hintzman, D.L. (1984) Minerva 2: A simulation of human memory. Behavior Research Methods, Instruments & Computers, 16, 96-101.
- Hunt, E.B., Marin, J., & Stone, P.J. (1966). Experiments in induction. Academic Press, New York, NY.
- Kolodner, J.L. (1986). Reconstructive memory: A computer model. Cognitive Science, 7, 281-328.
- Lockhart, R. S. (1969). Recency discrimination predicted from absolute lag judgments. Perception & Psychophysics, 6, 42-44.
- Luh, C.W. (1922). The conditions of retention. Psychological Monographs, 31, No 142.
- McGeoch, J.A. (1942). The psychology of human learning. New York: Longmans, Green and Company.
- McGeoch, J.A., & Irion, A.L. (1952). The psychology of human learning (2nd ed.), New York: McKay.
- Mensink, G.J., & Raaijmakers, J.G.W. (1988). A model of interference and forgetting. Psychological Review, 95, 434-455.
- Millward, R. (1964). Latency in a modified paired-associate learning experiment. Journal of Verbal Learning and Verbal Behavior, 3, 309-316.
- Murdock, B.B. (1982). A theory for the storage and retrieval of item and associative information. Psychological Review, 89, 609-626.
- Murdock, B.B. (1997). Context and mediators in a theory of distributed associative memory (TODAM 2). Psychological Review, 104, 839-862.
- Newell, A. (1990). Unified Theories of Cognition. Cambridge, MA: Harvard University Press.
- Postman, L., Kruesi, E., & Regan, J. (1975). Recognition and recall as measures of long-term retention. Quarterly Journal of Psychology, 27, 411-418.
- Raaijmakers, J.G.W., & Shrifin, R.M. (1981). SAM: A theory of probabilistic search of associative memory. In G.H. Bower (Ed.), The psychology of learning and motivation: Advances in research and theory (Vol 14). New York: Academic Press.
- Richman, H.B., Gobet, F., Staszewski, J.J., and Simon, H.A. (1996). Perceptual and memory processes in the acquisition of expertise: The EPAM Model. In K.A. Ericsson (Ed.), The road to expert performance: Empirical evidence form the arts and sciences, sports and games. Hillsdale NJ: Erlbaum Associates.

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.

Rizzuto, D.S., & Kahana, M.J. (2001). An autoassociative neural network model of paired-associate learning. Neural Computation, *13*, 2075-2092.

Rosenbloom, P. S., & Newell, A. (1987). Learning by chunking A production system model of practice. In R.S. Michalski, J. Carbonell, & T. Mitchell (Eds.), Machine Learning: An Artificial Approach II. Los Altos: CA: Morgan Kaufman.

Rundquist, W.N. (1965) Order of presentation and number of items as factors in paired-associate verbal learning. Journal of Verbal Learning and Verbal Behavior, *4*, 535-540.

Rundquist, W.N. (1966). Intralist interference as a function of list length and interstimulus similarity. Journal of Verbal Learning and Verbal Behavior, *5*, 7-13.

Sand, M.C. (1939). The effect of length of list upon retroactive inhibition when degree of learning is controlled. Arch. Psychol., N.Y., *133*, No 238.

Simon, H.A. (1966). A note on Jost's Law and Exponential Forgetting. Psychometrika, *31*, 417-418.

Simon, H.A. (1967). Motivational and emotional controls of cognition. Psychological Review, *74*, 29-39.

Simon, H.A. (1972). What is visual imagery? An information-processing interpretation. In L.W. Gregg (Ed.), Cognition in learning and memory. New York: Wiley..

Simon, H.A. (1976). The information storage system called "human memory." In M.R. Rosenzweig & E.L. Bennett (Eds.), Neural mechanisms of learning and memory. Cambridge, MA: The MIT Press.

Simon, H.A. (1979). Information processing models of cognition. Annual Review of Psychology, *30*, 363-396.

Simon, H.A. (1981). Otto Selz and information-processing psychology. In N. H. Frijda & A. deGroot (Eds.), Otto Selz: His contributions to psychology, The Hague: Mouton de Gruyter.

Simon, H. A., & Feigenbaum, E. A. (1964). An information-processing theory of some effects of similarity, familiarization, and meaningfulness in verbal learning. Journal of Verbal Learning and Verbal Behavior, *3*, 385-396.

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.

Suppes, P., Groen, G., & Schlag-Rey, M. (1966). A model for response latency in paired-associate learning. Journal of Mathematical Psychology, 3, 99-128.

Thune, L.E., & Underwood, B.J. (1943). Retroactive inhibition as a function of degree of interpolated learning. Journal of Experimental Psychology, 32, 185-200.

Underwood, B.J. (1953). Studies of distributed practice: VIII. Learning and retention of paired non-sense syllables as a function of intralist similarity. Journal of Experimental Psychology, 45, 133-142.

Underwood, B.J. (1982). Paired associate learning: Data on pair difficulty and variables that influence difficulty. Memory and Cognition, 10, 610-617.

Voss, J.F., Thompson, C.P., & Keegan, J.H. (1959). Acquisition of probabilistic paired-associates as a function of $S-R_1$, $S-R_2$ probability. Journal of Experimental Psychology, 58, 390-399.

Weinstock, R.B. (1976). Recognition of paired associates is not necessarily contingent on stimulus recognition. Psychological Reports, 38, 643-650.

Wynn, W. H. (1966). An information-processing model of certain aspects of paired-associate learning. Unpublished doctoral dissertation, University of Texas, Austin, TX.

Yntema, D. B., & Trask, F. P. (1963). Recall as a search process. Journal of Verbal Learning and Verbal Behavior, 2, 67-74.

Zhang, G., & Simon, H.A. (1985). STM capacity for Chinese words and idioms: Chunking and the acoustical loop hypothesis. Memory and Cognition, 13, 202-207.

Appendix

Routines used by EPAM for Responding and Learning

There are two overall routines used by EPAM when responding in the paired-associate paradigm (*Recognize* and *Find-associate*) and one overall learning routine (*Study-paired-associate*). *Recognize* is used by both of the other routines so we will discuss it first. The *Study-paired-associate* routine makes use of EPAM's various methods of learning which include *Study* which grows the discrimination net, *Familiarize* which adds information to an entry and updates its time tag, and *Create-link-to* which insures that discrimination net nodes point to entries

Recognize Routine

The *Recognize* routine is the one called when the simulated subject is asked to answer the question "Do I recognize the stimulus response pair that was presented?" It is given two main inputs, the stimulus and the response and it outputs an entry, if one can be found, that matches the stimulus and response.

The recognize routine also has another major input: the type of entry desired which functions as "directed association" theorized by Otto Selz (Simon 1981). As a result, EPAM can answer different questions of the stimulus such as "What *episode* is associated with this stimulus?" or "What *category* is associated with this stimulus?" In the paired associate simulations of this paper, the entries being created all have "episode" as their type, while in our classification learning simulations, the entries all have "category" as their type.

The first step of the recognize routine is to create a stimulus-response object to be recognized. This stimulus response object will unite the stimulus and response so that EPAM can search its memory to find a match. For example if the stimulus is S_1 and the response is R_1 then the stimulus-response object would be the list $(S_1 R_1)$.¹⁵

There are two versions of recognize. The "quick" version and the standard version. The "quick" version is used in the fan-effect experiments because the subject has been instructed to respond as quickly as possible. The standard version is used when the subject has more time. The difference between two is that the quick version only tries to find the entry at either the stimulus or the response (50-50 chance of each) while the standard version searches both the stimulus and response (in a random order).

At whichever location (stimulus or response) that is chosen for investigation, *Recognize* calls a routine

called *Find-correct-episode-attached?* in order to determine whether the stimulus-response object matches an entry that is pointed to at the node to which the stimulus or response sorts. If *Find-correct-episode-attached?* exits with an entry, then *Recognize* outputs the same entry. If *Find-correct-episode-attached?* exits “nil” then the *Quick-recognize* routine exits “nil,” but the standard version tries the other location. (If *Recognize* had first tried the stimulus location, it would now try *Find-correct-episode-attached?* at the response location and vice versa.)

Find-correct-episode-attached? Routine

The *Find-correct-episode-attached?* routine is given three inputs: (1) the node to be investigated, either the node to which the stimulus sorted or the one to which the response sorted; (2) the stimulus-response-object, and (3) the type of entry desired, in these experiments the type is “episode.” Its output is either an entry that matches or nil. It follows the following steps:

1. It first finds the list of pointers to entries of the particular type (in this case the “episodes”) that are associated with the node to be investigated (i.e. the node to which the stimulus sorted or the node to which the response sorted).
2. Next, it generates the episodes in perceived recency order until one is found that matches the stimulus-response object or until all of the episodes have been generated without a match. It calls the routine *Not-different?* to perform the match between the stimulus-response object and the entry.
3. It charges the system an additional 250 msec. for each episode after the first that was generated in Step 2.
4. It exits with the first episode found that is *Not-different?* from the stimulus-response object or exits with nil if all of the episodes were generated with no match. A pointer to the episode found is placed in the “focus cell” of declarative memory called the “*RSF*.”¹⁶

Not-different? Routine

The *Not-different?* routine compares an object with an image and exits “t” if the two match, or “nil” if the two do not. It is a recursive routine which calls itself to examine subobjects and subimages. When called by *Find-correct-episode-attached?* the object that is input is the stimulus-response object and the image is the image of the entry that was generated. An object and image do not match if the image contains any value

or subobject that is not on the object

Assume, for example, that the stimulus-response object is (XLC, 1) and the image is (X—, 1). Then the image and object would match because the image does not contain any subobjects that are not on the object. On the other hand if the stimulus-response object is (XLC, 1) and the image is (XF-,1) or (X—, 2) then the object and image would not match.¹⁷

Find-associate Routine

The *Find-associate* routine is given two required inputs, the stimulus and the type of entry, and an optional input, the position to be output. In most of the experiments simulated in this paper, the position to be output is the second position, however in the backward-recall simulation, the position is the first position since in that experiment, the subject is given a response and asked to respond with the stimulus.¹⁸ The output of *Find-associate* is the response if one is found, or “nil” if no response is found.

Find-associate creates a stimulus-response object for use by the *Recognize* routine. Normally, when *Recognize* is called, the stimulus response object given to it would be (S₁ R₁), but the stimulus-response object created by *Find-associate* is (S₁ _). The underscore indicates that the response is unknown. In the backward-recall simulation, the stimulus-response object that is created is (_ R₁).

The actual work of the find-associate routine is done by the *Find-response-at-stimulus-chunk* routine which calls *Recognize* with the stimulus-response object (S₁ _) and then responds with the response part of the matched entry. For example, if the stimulus-response object is (XLC, _) and the entries at the stimulus node were those shown in Figure 3b and were generated for *Recognize* in the perceived recency order (XQ-.4), (XF-,3), (X—,1). Then the first two would be rejected, and the third (X—,1) would match the entry, so the output would be the response “1.”

Study-paired-association Routine

After responding to a stimulus, EPAM calls the *Study-paired-association* routine. The inputs are the type of entry (“episode” in all of these simulations) the stimulus object, the correct response-object and whether the response that was just made by the simulated subject was correct. EPAM does the following steps and exits after the first step succeeds.

1. If EPAM learning is busy as indicated by an variable called the *epam-clock* which tells when current EPAM learning will be complete, then EPAM will exit without doing any learning.
2. If EPAM’s response was correct, then 30% of the time (this percentage is set by the *SAL-C*

parameter) EPAM calls *Study-when-overlearning*. The other 70% of the time, EPAM will exit without any learning.

3. When there is more than one episode associated with a stimulus, then 50% of the time (this percentage is set by the *SAL-A* parameter) EPAM calls *Study* with the stimulus in order to discriminate the stimulus and then calls *Associate* to create or improve the entry that associates the stimulus-object with the correct-response object.
4. In all other situations, EPAM simply calls *Associate*, which usually creates a new entry or just updates the time tag of an existing entry.

Study-when-overlearning Routine

This routine calls *Study*, 50% of the time, to further discriminate the stimulus or calls *Study*, the other 50% of the time, to further discriminate the response. Then it calls *Associate* to refresh the association between the stimulus and the response. (When it calls *Associate* it specifies as the value of “correct?” either “stimulus” or “response” depending upon which one was just studied.)

Study Routine

Study is EPAM’s discrimination net growing routine. In these experiments tests are always added in the order of positions on the list. Thus a test for first position is the first added, then a test for second position is the next. Whenever possible, *Study* adds a new branch to the discrimination net to further discriminate the object that is being studied. Normally it takes EPAM 4 seconds to add a new node to the discrimination net, but if a test has to be added it takes EPAM 12 seconds (8 seconds for the test and 4 seconds to build the new node).

Associate Routine

The *Associate* routine is the routine used by EPAM for creating or improving an entry that unites a stimulus with a response. It always calls *Familiarize* which, among things, updates the time-tag of the entry. It also always calls *Create-link-to* in order to insure that pointers to the entry reside at the stimulus node and at the response node. Its required inputs are the type of entry (“episode” in these simulations), the stimulus, and the response.¹⁹ The optional input “correct?” is true if *Associate* was called as part of overlearning after the system had just responded with the correct response.

The *Associate* routine, once called, is assumed to operate in the background. The simulated subject

can be rehearsing the stimulus and response or even respond to another stimulus while *Associate* completes its work..

Associate does one or other of several different steps depending upon the current status of the association:

1. If there is no subnet for the current “is-a” of the stimulus or response, then it creates a subnet but continues to the next step. All of the stimuli and responses in these simulations are part of a subnet called “syllable.”
2. *Associate* then creates a stimulus-response object linking the stimulus and response ($S_1 R_1$) and then calls *Recognize* which searches memory for an entry that matches the stimulus-response object. If one is found, then *Recognize* places the appropriate entry in the focus cell of declarative memory. If no entry can be found that matches, then *Associate* creates a new entry in the focus cell. In the slotted-schema simulation, the new entry that is created was two slots, indicated by underscores (), but in most simulations, a new entry starts as an empty list ().
3. If this is an overlearning situation and the stimulus was just studied then *Familiarize* is called at the surface level to update the stimulus part of the entry and *Create-link-to* is called from the stimulus to the entry so that a pointer will be created from the new stimulus node to the entry. Also *Create-link-to* is called to add a pointer (if not already there) from the response node to the entry. Then normally *Associate* exits. However, if *Study* had not succeeded with moving the stimulus to a new node, *Move-episode-up-on-ladder* is called before *Associate* exits.
4. If this is an overlearning situation and the response was just studied then *Familiarize* is called at the surface level to update the response part of the entry and *Create-link-to* is called from the response to the entry so that a link will be created from the new response node to the entry. Then *Associate* exits.
5. In any other event the *Associate-if-error* routine is called, after which *Create-link-to* is called and then *Associate* exits.

Familiarize

Familiarize is the routine that is used to add information to an entry of declarative memory. Its inputs are usually a stimulus-response object and a declarative-memory entry. Other inputs can direct *Familiarize*

to add information to a particular part of the entry or can direct *Familiarize* to just add information at the surface level (which would be just enough additional information to sort to the stimulus and to the response).

Each call to *Familiarize* can add some information from the stimulus-response object to the entry. The learning time is 4 seconds per chunk of information added. However, if the information can be placed in preexisting slots then the time taken to add information is just 250 ms. per slot. The *Familiarize* routine also automatically updates the time-tag at the entry.

Create-link-to Routine

The *Create-link-to* routine checks to see if there is already a pointer from a discrimination net node to an entry. If there is not, then it adds such a pointer with a learning-time cost of 1 second.

Move-episode-up-on-ladder

When there are multiple entries at the stimulus node of this type of entry (i.e. if there are multiple “episodes” at the stimulus node and if the entry in question is not the original episode at the node, then pop the oldest episode.) The result of this action will mean that the entry being studied will be more stable and could possibly become the permanent entry of this type at this node. When the oldest entry of this type is popped from the node, the time cost is 4 seconds.

Associate-if-error

The inputs to this routine are the stimulus-response object and the position of the response on the stimulus-response object. The following steps are taken with this routine exiting immediately if one is taken.

1. If the response part of the entry is already complete, then *Familiarize* the entry at the surface level and exit.
2. If there is no response on the entry, then *Study* the response, and repetitively *Familiarize* the entry at the surface level so that the information necessary to sort to both the stimulus and to the response is included in the entry.
3. If the response is not complete on the entry, then *Familiarize* the response part of the entry then *Familiarize* the entry at the surface level so that the information necessary to sort to both the stimulus and to the response is included in the entry.
4. If none of the above apply, then just *Familiarize* the entry at the surface level.

End Notes

¹ The four simulations simulated with the full model are: (1) the *constant learning time* simulation, (2) the *prefamiliarization of stimuli and responses* simulation, (3) the *effects of rehearsal strategy* simulation, and (4) the *slotted-schema* simulation.

² We have adopted a parameter in the EPAM VI model which functions as a sort of combination of “a” and “b” of the SAL simulations. Named “*SAL-A*,” it is set at 50% in all EPAM VI simulations. If, after an error, there are already multiple entries at a stimulus, 50% of the time EPAM will study the stimulus before studying the association. The other 50% of the time EPAM will simply study the association. EPAM doesn’t need a separate parameter to slow down the learning process in these simulations. The time required to add information to long-term memory severely limits learning speed in this paradigm.

³ Hintzman simulated four additional experiments with SAL I. He found that SAL I successfully simulated the effects of *probability matching* and *massed vs. spaced* item repetition found by Greeno (1964) but failed to simulate the effects of *whole vs. part* learning found by McGeoch and Irion (1952) or the effects of list length found by Calfee and Atkinson (1965), Runquist (1965, 1966) and Sand (1939). We have not yet taken the time to simulate the effects of *massed vs. spaced* item repetition or those of *whole vs. part* learning. We have, however, simulated the effects of list-length with the same disappointing results found with SAL. Although people make more mistakes on average per item when learning long lists as compared to short lists, neither EPAM VI nor SAL do so (unless similarity is confounded with list length). Perhaps there is a “discouragement” factor involved, not yet captured in EPAM VI, which makes it harder for humans to learn longer lists. Simon (1967) discusses how discouragement could be programmed into an information-processing model. Our simulation of probability matching is discussed later in this paper.

⁴ These were the same stimuli used with SAL I, but slightly different from the stimuli that Hintzman (1969) used with his human subjects.

⁵ Actually, only one test node would have been sufficient to discriminate the low-similarity stimuli from each other. EPAM VI creates the additional nodes as part of the overlearning process that we will discuss later.

⁶ Steps 4 through 7 usually occur together during a single call to the long-term learning process.

⁷ The very gradual rise in EPAM VI’s total learning time occurs because we have programmed overlearning to be independent of presentation rate. Overlearning is attempted on 30% of the presentations, with no allowance for length of presentations. As a result it is called less frequently during the experiments that have slower presentation times, and thus fewer trials.

⁸ The ratio is higher for EPAM VI than EPAM III partly because EPAM VI spends time over-learning while EPAM III does not.

⁹ The full version was needed for the simulation of the constant learning time simulation so that EPAM VI would take advantage of presentation times as long as 19 seconds. It was also needed in the simulation of the “one-at-a-time” strategy in the effects of rehearsal strategy simulation. All of the other simulations could just as easily be simulated with either version.

¹⁰ There is another task also called “recognition” in which subjects are simply asked to make an “old-new” recognition rating about whether or not they have studied a stimulus or a response. Performance on this task appears to be based upon a different mechanism (Underwood, 1982; Weinstock, 1976).

¹¹ We are not saying that people always use these time tags to make recency judgments. Sometimes people may have actual dates associated with their memories such as, “that was just before our son’s fifth birthday.”

¹² We estimated Lockhart’s data (shown as “X’s” on the graph) by using a ruler and measuring distances on his published graphs. In order to draw the curve, we ran EPAM VI 1,000,000 times at each point, using computer-generated random numbers and consulting a table of normal curve equivalents that had been programmed into EPAM.

¹³ We use 250 ms. in all of our EPAM simulations as the time required to do a conscious act. This is the average amount of time required by a reader to do a single eye fixation and also the amount of time taken in our simulation for an expert mnemonist to move his focus from one node to another in his retrieval structure (Richman et al. 1995).

¹⁴ Bower's (1972) human subjects were instructed to associate the concrete nouns by imaging a visual scene or mental picture in which they are interacting in some way, a task which Simon (1972) analyzed as having the same result as the meaningful-sentence task.

¹⁵ Actually, the stimulus-response object and the image of the entry would have the form $((\text{is-a episode}) S_1 R_1)$ but we have left out the "is-a" of the entry for simplicity of presentation.

¹⁶ The focus cell in declarative memory is called "*RSF*" because in EPAM VI it was the "Retrieval Structure Focus." When EPAM IV simulated traversals of the retrieval structure in its simulation of expert memory (Richman et al., 1995) pointers to nodes of the retrieval structure were moved in and out of the focus cell with a time cost of 250 ms. to change the pointer.

¹⁷ The dashes indicate missing information.

¹⁸ There are additional optional inputs which do not figure in these simulations. One is "instant" which is true when find-associate is to respond with the first entry without checking to make sure that the entry matches the stimulus and "rsf" which is true in a serial learning experiment where find-associate is used to find the first, last, or next item on the entry which is already in the focus cell of declarative memory.

¹⁹ Its optional inputs position1 and position2 allow the location to be used as the stimulus and associated with the person in the fan-effect experiments even though EPAM will put the two in (person, location) order on the entry. The optional input "rsf" is true in serial learning experiments where associate is used to put a response in the first or last position, or before or after another object on the entry which is already in the focus cell of declarative memory.