

Psychological Review

VOLUME 88 NUMBER 2 MARCH 1981

Search of Associative Memory

Jeroen G. W. Raaijmakers
University of Nijmegen
Nijmegen, The Netherlands

Richard M. Shiffrin
Indiana University

A general theory of retrieval from long-term memory combines features of associative network models and random search models. It posits cue-dependent probabilistic sampling and recovery from an associative network, but the network is specified as a *retrieval structure* rather than a storage structure. The theory is labeled SAM, meaning *Search of Associative Memory*. A quantitative simulation of SAM is developed and applied to the part-list cuing paradigm. When free recall of a list of words is cued by a random subset of words from that list, the probability of recalling one of the remaining words is less than if no cues are provided at all. SAM predicts this effect in all its variations by making extensive use of interword associations in retrieval, a process that previous theorizing has dismissed.

A cue-dependent probabilistic search theory of retrieval has been developed to operate within a retrieval structure based on an associative network. This theory, referred to as SAM (Search of Associative Memory), has been applied to the results of studies investigating free recall, paired-associate recall, and recognition. The theory is capable of fitting the results simultaneously, with no essential change in processes or parameters. Raaijmakers (1979) and Raaijmakers and Shiffrin (1980) give early results supporting these contentions.

In the present article, our goal is the delineation of the general theory and a thor-

ough exposition of an application to one particular research setting. For expository purposes, we begin with a simplified model and application and then turn to the general theory. The first section of the article describes a relatively simple computer simulation program for free recall that serves to illustrate the main features of the theory. The power of the approach is then illustrated by applying the simulation model to a puzzling phenomenon known as the part-list cuing effect (e.g., Slamecka, 1968). Section II of the paper discusses in detail research and theory directed toward the part-list cuing effect; Section III explores in depth the way in which the simulation model deals with this phenomenon. It should be noted that the model presented in Sections I and III is a precisely stated, carefully delineated, scaled-down version of the general theory aimed at a limited set of paradigms. Discussion of the rationale behind many of the assumptions, the ways in which the theory handles other types of paradigms, the strengths and weaknesses of the approach, and comparisons with other theories are reserved for Section IV.

This research was supported in part by a fellowship from the Netherlands Organization for the Advancement of Pure Research to the first author and Public Health Service Grant 12717 and National Science Foundation Grant BNS-77-00156 to the second author. We gratefully acknowledge the help of Henry Roediger III, Edward Roskam, Susan Dumais, David Foyle, Bruce Williams, and Gary Gillund, who assisted us in the various stages of development of this work.

Requests for reprints should be sent to R. M. Shiffrin, Psychology Department, Indiana University, Bloomington, Indiana 47405.

I. A Simulation Model for Free and Cued Recall

The general theory will be abbreviated as SAM (Search of Associative Memory); the simulation model for free recall will be abbreviated as SAMS (SAM Simulation).

The Paradigms

SAMS is developed to deal with the paradigms of free verbal recall and free recall with added cues. In free recall, a list of n words is presented in random order at a rate of t sec per word. Both t and n may be varied between lists. In immediate free recall, the subject is asked to recall as many words as possible, in any order, immediately following list presentation. In this case, the most recently presented words are recalled at an enhanced level (the *recency effect*), presumably reflecting the presence of those words in a short-term store. In delayed free recall, an interpolated task, usually arithmetic, is interposed between presentation and test. The interpolated task is presumed to clear short-term store of the list words so that all recall is from long-term store. Indeed, the recency effect is eliminated in this case. The part-list cuing paradigm will be described in detail in Section II. In brief, at the time of testing, some of the list words are presented to the subject, supposedly as helpful aids to enable recall of the remaining list words.

The data from such experiments are typically analyzed in terms of the number of recalled words or in terms of the probability of recall (which is sometimes analyzed in terms of the presentation position or the output position of particular words). In addition, the rate of output is often analyzed, sometimes in terms of interresponse times, but more often in terms of the cumulative number of items recalled as a function of output time.

Short-Term Store and Long-Term Store

SAMS utilizes a two-phase memory system (Atkinson & Shiffrin, 1968; Shiffrin, 1975). Short-term store (STS) is the temporary store into which information about presented items is placed and in which con-

trol processes such as coding and rehearsal are carried out. Long-term store (LTS) is the permanent store, containing all prior information plus new information transferred from STS. Retrieval from LTS is quite fallible, as described below.

STS is limited in capacity (Shiffrin, 1976), so that only a limited number of items may be retained, rehearsed, and coded at one time. In particular, SAMS assumes a buffer rehearsal system operating as follows: Each presented word enters STS, joining previous words there, until the buffer size, r , is reached. Then each new word entering the buffer replaces one of the r words already present. The word to be replaced is chosen with probability $1/r$.

If immediate testing is used, the r items currently in the buffer are all correctly recalled, and then retrieval from LTS begins. If delayed testing is used, it is assumed that the arithmetic clears the buffer at the same rate that would have obtained had the list continued. After a long period of arithmetic, the buffer will be empty of list items at test, and retrieval will begin from LTS at once.

Storage in Long-Term Store

What is stored in LTS are associative relationships between context and word information, and between a set of context-plus-word information for one word and a set of context-plus-word information for another word. The former is termed *word-context information*, and the latter is termed *word-word information*. The amount of word-context information stored is assumed to be proportional to the total amount of time a given word remains in the rehearsal buffer. The amount of word-word information stored is assumed to be proportional to the total amount of time those two words are simultaneously present in the rehearsal buffer.

Word information is information that enables the subject to produce the name of the encoded word. Context information refers to temporal and situational factors that might also be present in STS at a given time, including environmental details, physical sensations, emotional feelings, and all thought processes not directly relevant to name production.

Besides the storage that takes place during list presentation, additional storage may occur during the long-term retrieval process itself. The assumptions governing such storage are given below.

Long-Term Retrieval

LTS retrieval is assumed to be a cue-dependent process operating on a retrieval structure based on an associative memory. At each step of the search process, cues are assembled in STS and used as a probe of LTS. A localized set of information in LTS, called an *image*, is *sampled*. The sampling probabilities are a function of the strength of association between the probe cues and the various images in LTS. The information in the sampled image is then accessed and evaluated in a process called *recovery*. The probability that enough information is recovered to enable the name of the encoded word to be recalled is a function of the strength of association between the probe cues and the sampled image.

In SAMS, the only cues considered are the general context cues, called C_T , and the words from the presented list, termed $W_{1T}, W_{2T}, \dots, W_{nT}$. The subscript T indicates the word at test. A probe set always consists of C_T alone or of C_T along with one of the word cues, for example, (C_T, W_{iT}) . The images that can be sampled consist of the stored information about each of the n words in the presentation list, including the context at that time; these are denoted $W_{1S}, W_{2S}, \dots, W_{nS}$. The subscript S indicates the information as it is stored.

The Retrieval Structure

SAMS assumes that relationships between the probe cues and the images may be completely summarized in a *retrieval structure* that gives the strengths of relationship between each possible probe cue and each possible image. This is illustrated in the matrix of Figure 1. In the figure, $S(C_T, W_{iS})$ represents the strength with which the context cue tends to sample the image of Word i ; $S(W_{jT}, W_{iS})$ represents the strength with which Cue-word j tends to sample the image of Word i .

		IMAGES	
		W_{1S}	W_{nS}
CUES	C_T	$S(C_T, W_{1S})$	$S(C_T, W_{nS})$
	W_{1T}	$S(W_{1T}, W_{1S})$	$S(W_{1T}, W_{nS})$
	\vdots		
	W_{nT}	$S(W_{nT}, W_{1S})$	$S(W_{nT}, W_{nS})$

Figure 1. The *strength matrix*—the matrix of strengths that determine the probabilities of selection and recovery of list images (horizontal margin) when various cues (vertical margin) are used in the cue set. (Entries in the cells are strengths from individual cues to individual images; when multiple cues are used in the cue set, then the strengths are combined according to Equations 1 and 2 in the text. C_T refers to the context cue; W_{iT} and W_{iS} refer to the Cue Word i and the image of Word i .)

The strengths in this matrix are assumed in SAMS to be proportional to the amounts of information stored during presentation. Let t_i be the time that Word i stays in the buffer, and let t_{ij} be the time that Words i and j are together in the buffer simultaneously. Then

$$S(C_T, W_{iS}) = at_i; \quad S(W_{iT}, W_{jS}) = bt_{ij};$$

$$S(W_{iT}, W_{iS}) = ct_i,$$

where a , b , and c are parameters to be estimated.

It should be noted that many pairs of words will never occupy the buffer simultaneously. In all such cases, a small residual strength value is assumed. In particular, if $t_{ij} = 0$, $S(W_{iT}, W_{jS}) = d$, where d is another parameter to be estimated (but is presumably smaller than a , b , or c).

The Sampling Rule

Equations 1a and 1b give the quantitative sampling rules:

$$P_S(W_{iS}|C_T) = \frac{S(C_T, W_{iS})}{\sum_{j=1}^n S(C_T, W_{jS})}, \quad (1a)$$

$$P_s(W_{is}|C_T, W_{kT}) = \frac{S(C_T, W_{is})S(W_{kT}, W_{is})}{\sum_{j=1}^n S(C_T, W_{js})S(W_{kT}, W_{js})} \quad (1b)$$

In Equation 1, P_s indicates the sampling probability. The left-hand sides of these equations give the probability of sampling the image of Word i , given that the probe cues consist of C_T in Equation 1a, or C_T and W_{kT} in Equation 1b. Equation 1a is a simple ratio rule for strengths. Equation 1b is a ratio rule for the products of strengths. This product rule has the effect of focusing the search so that there is a tendency to sample images with high strengths to both cues, rather than images with high strengths to only one of the cues.

It should be noted that when at_i is substituted for the strengths in Equation 1a, the a cancels, and the rule becomes a ratio of rehearsal times. However, when similar substitutions are made in Equation 1b, the parameters a , b , c , and d do not cancel (except for special combinations of values).

The Recovery Rule

Equations 2a and 2b give the probability that the subject will be able to generate the name of the word encoded in Image i , given the cue (or cues) that made up the probe set during sampling, and given that this recovery instance is the first for Image i with that probe set. In Equation 2, P_R indicates the recovery probability.

$$P_R(W_i|C_T) = 1 - \exp\{-S(C_T, W_{is})\} \quad (2a)$$

$$P_R(W_i|C_T, W_{kT}) = 1 - \exp\{-S(C_T, W_{is}) - S(W_{kT}, W_{is})\} \quad (2b)$$

These equations essentially say that recovery probabilities rise as the cue-to-image strengths rise. Note that the parameters a , b , and c do not cancel out of these recovery equations; in fact, these parameters may be viewed as recovery parameters.

SAMS makes special assumptions when an image that has already been sampled with a given probe set is again sampled with that same probe set later in the search. These

"conditionalization" assumptions hold that the outcome of the second and subsequent recovery attempts will match the outcome of the first recovery attempt. In addition, if the context cue and any other word cue together do not lead to recovery of Image i , then it is assumed that any subsequent recovery attempt for Image i with the context cue only will also fail. Thus a new independent recovery chance occurs whenever the probe set contains a new cue (i.e., one not utilized before for the image sampled). These assumptions represent an assumption that successive recovery attempts of the same image with the same probe set are based on essentially the same information.

Storage During Retrieval

Whenever a successful recovery occurs, the strengths of the probe set to the recovered image are increased. This process is termed *incrementing* and obeys the following rules:

$$S'(C_T, W_{is}) = S(C_T, W_{is}) + e;$$

$$S'(W_{iT}, W_{js}) = S'(W_{jT}, W_{is}) \\ = S(W_{iT}, W_{js}) + f;$$

$$S'(W_{iT}, W_{is}) = S(W_{iT}, W_{is}) + g. \quad (3)$$

Here the primes represent the strengths after incrementing, and e , f , and g are the incrementing parameters. Two special assumptions are made concerning incrementing: First, any increment of a context to word strength is always linked to an increment of the recovered image's self strength. That is, the e and g increments always occur together. Second, only one increment can ever take place for any given cue-to-image strength. (See Section IV, *Incrementing*, for further discussion.)

The Search Process

The retrieval process involves an extended search that requires a series of decisions on the part of the subject concerning what cues to use during each sampling phase and when to stop searching. The retrieval process in an arbitrary setting is diagrammed in Figure 2. The subject makes a plan, chooses a set

of probe cues, samples and tries to recover, decides whether to stop searching, and then loops back to the retrieval plan to begin the next phase. This general system will be discussed in Section IV. The specific decisions and control processes used in SAMS are given in the flow chart shown in Figure 3.

The stopping rule. A failure is defined as any sampling and recovery attempt that does not result in recall of a new word. The main phase of the retrieval process is assumed to end whenever the cumulative total of failures in the search reaches a value termed K_{MAX} .

Selection of probe cues. If cues are not provided, the subject begins by using context only as a cue. This probe set continues to be used until some new word is recalled. Whenever a new word is recalled, the next probe set consists of two cues: the context cue and the word just recalled. This rule applies even when the previous probe set consisted of context plus a different word.

Stopping rule for word cues. L_{MAX} is the stopping criterion for a word cue. If a given combination of context plus word serves as a probe set for L_{MAX} consecutive loops of the search, and no new word is recalled, then the subject drops the word from the probe set and uses context alone.

Rechecking. When the main search phase ends (when K_{MAX} is reached), then *rechecking* begins. In this phase, each previously recalled word (either from the main phase or rechecking) is used along with context in a probe set. Each such probe set is utilized until L_{MAX} failures accumulate. Any new words recalled during this period are saved and are used as cues during later rechecking. The rationale for the rechecking process is based on a certain arbitrariness in the assumptions concerning selection of cues for probe sets. A word may be recalled and used as a cue. Then it may happen that a new word is recalled on the very next loop of the search. Our rules require that the probe set be changed to use this new word, but it could be argued that the previous word cue is still of value and has not been "used up." The incorporation of a rechecking process allows us to argue that all associative retrieval routes have been checked thoroughly, at least up to an L_{MAX} stopping criterion.

RETRIEVAL FROM LONG-TERM-STORE

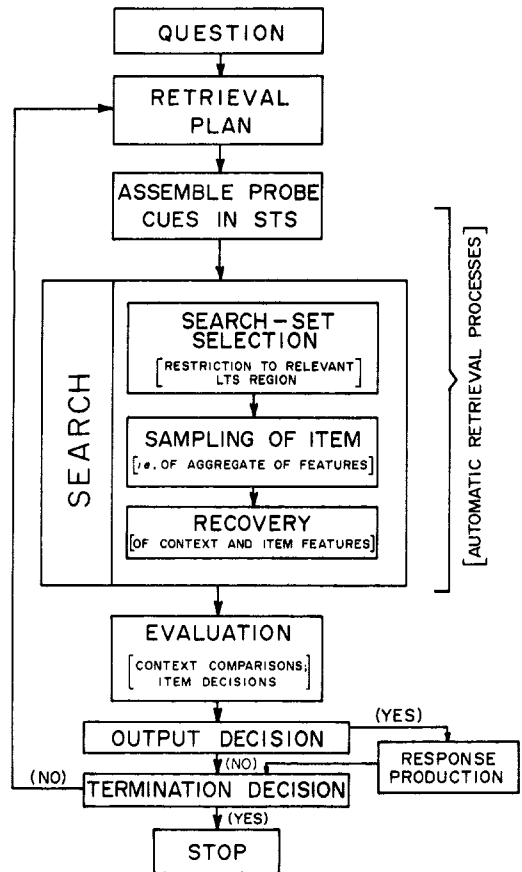


Figure 2. A generalized depiction of the various phases of retrieval in the theory. (STS = short-term search; LTS = long-term search.)

Parameters

The SAMS model described above is ready to be applied to the data from free-recall studies. The parameters are a (context-cue-to-image strength), b (word-cue-to-image strength), c (word-cue-to-self-image strength), d (residual-word-cue-to-image strength), e (context-to-image increment), f (word-cue-to-image increment), g (word-cue-to-self-image increment), K_{MAX} (total failure stopping criterion), L_{MAX} (stopping criterion for a word cue), and r (buffer size).

At first glance, 10 parameters seems quite a high number to fit the results from free-recall studies, especially since Shiffrin (1970)

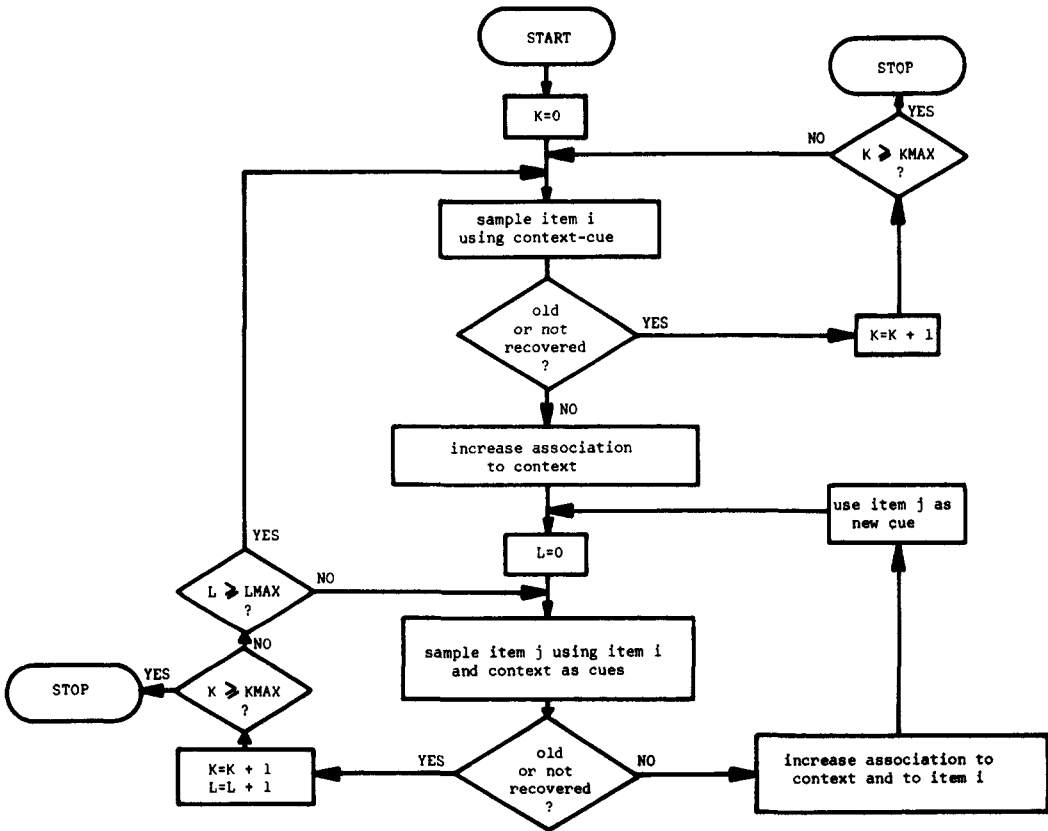


Figure 3. A flowchart for Phase 1 of the retrieval process in the computer simulation of SAMS developed for free recall. (K and L are counters of failures. *Old or not recovered* refers to a failure to recall a new item.)

used a much simpler search model with just three parameters to fit a great deal of free-recall data. This objection is ameliorated by the following factors. We can show that most of the present parameters, and their precise values, are not essential for the fit of the model to most of the data. The parameters are listed above for generality, even though some are never varied and others are equated before fits to the data are begun. Some of the parameters are given nonzero values and included in the fit merely to demonstrate that the presence of the processes they represent will not harm the ability of the model to predict the data. In fact, we have set many of these parameters to zero, and no harm to the model's predictions has resulted. However, each of these parameters represents processes that we feel are needed on logical

grounds, or needed to deal with data that will not be discussed in detail in this article (see Raaijmakers & Shiffrin, 1980). The roles played by the various parameters have been extensively explored by simulation means, as have certain process assumptions, and the results of these explorations will be either summarized briefly or reported in detail in the remainder of the article.

Free-Recall Data

In this section we will apply SAMS to free-recall data in order to get a set of reasonable parameter estimates. A set of data particularly well suited for the initial explorations of the model was collected by Roberts (1972). Four list lengths (10, 20, 30, or 40 words) and five presentation rates (.5, 1,

2, 4, or 8 sec per word) were covaried. No interpolated arithmetic task was used (so both short- and long-term components are assumed in recall). Both visual and auditory presentation modes were utilized, but we have averaged the results for these two modes. Our main concern is average recall performance for each type of list (rather than performance as a function of serial presentation position).

The results are given in the upper panels of Figures 4 and 5. Clearly, mean number of words recalled is a negatively accelerated

function of total presentation time for a fixed list length (Figure 4; see also Waugh, 1967). Also, equal total presentation times for different list lengths do not yield equal recall; rather, longer lists yield more total words recalled (Figure 4). Finally, as list length increases for fixed presentation time per item, total recall increases almost linearly, though with a slight hint of negatively accelerated growth (Figure 5). The "list length effect" is also implicit in these data: The probability of recall of a given word decreases as list length increases.




Figure 4. Observed data (Roberts, 1972; top panel) and predictions of SAM (lower panel) for mean words free recalled as a function of presentation time and list length (LL). (The parameters of the model are $r = 4$; $K_{MAX} = 30$; $L_{MAX} = 3$; $a = c = .10$; $b = .10$; $d = .02$; $e = f = g = .70$. These parameters are used to generate predictions in all the following figures, except where indicated differently.)




Figure 5. The same data are presented as in Figure 4, except that the label of each curve is presentation time (PT) per item.

Parameter Estimation

The model has characteristics that make analytical derivations infeasible; it is therefore constructed as a computer simulation program. Predictions are derived by Monte Carlo methods—by averaging the results of large numbers of “statistical subjects.” Parameter estimates are obtained by generating predictions for various combinations of parameter values and choosing the best combination. The number of simulations (NSIM) was 300 for the predictions shown in the figures. Nothing like an exhaustive search was used to obtain parameter estimates: Some parameters were set on the basis of past data

and models, some were arbitrarily set equal to others, and some were varied over small regions of values. Surely, then, the fits are not optimal. They are, however, quite satisfactory for our purposes. The parameter values, methods of estimation, and some general comments are discussed next.

Parameter r. The buffer size was set equal to 4 on the basis of previous fits to serial position curves and on the basis of numerous estimates of short-term capacity as 2.5 (see Crowder, 1976). Typical methods of estimating short-term capacity, such as that of Waugh and Norman (1965), produce values much lower than the buffer size; it can be shown that our predictions, using a

buffer size of 4, when analyzed according to the Waugh and Norman procedure, result in an estimate of short-term capacity of 2.5. It should be mentioned also that our model gives an excellent account of the Murdock (1962) serial position data when r is set equal to 4.

Value of K_{MAX} . The criterion for cessation of search was set equal to 30 failures. This value was chosen in part because Shiffrin (1970) fit Murdock's data with an assumption of about 36 samples (corresponding to about 30 failures). It can be shown that the choice of stopping rule and the value of 30 are not crucial for predicting the qualitative and quantitative effects in the present data (nor in much other data, as well). Of course, if the value of K_{MAX} is raised, predicted recall will increase slightly, but the increase may be counterbalanced by small changes in the other parameters.

In general, K_{MAX} is a decision criterion, whose value may be chosen by the subject and manipulated by instructions.

Value of L_{MAX} . The number of failures of an item cue before it is discarded and context only is utilized as a cue was set equal to 3. The performance changes expected as a result of changes in L_{MAX} are fairly difficult to ascertain, especially since increases in L_{MAX} extend the period of rechecking. With the help of simulations, the following general properties have been determined: When interitem strength attains values near or below those estimated for Roberts' (1972) data, then a word cue's tendency to sample itself is fairly high, and a context cue alone can lead to more efficient sampling. (This reasoning does not hold if interitem strength or residual strength is raised.) Since large values of L_{MAX} reduce both the number of context cues utilized in the search, and also the number of different word cues utilized, increases in L_{MAX} tend to decrease performance. This effect is counterbalanced somewhat by the fact that recovery probability for word-plus-context cue sets is greater than for context-only cue sets. The net result is that, without rechecking, increases in L_{MAX} from 1 to 5 have only a small effect, a slight decrease in performance. However, the amount of search during rechecking increases as L_{MAX} increases, so in the model

as it stands, performance turns out to be a weakly inverted-U-shaped function of L_{MAX} , with a maximum near 3. Thus we fortuitously chose a value of L_{MAX} that maximizes performance. Also, it certainly is the case that extensive use of interword associations is made during the search when $L_{MAX} = 3$, a desirable property of the model.

The value of L_{MAX} is a choice of the subject and may be manipulated by instruction.

Parameter a. The context-to-image strength per second of rehearsal time was one of the three estimated parameters and attained a value of .10. Sampling probabilities are not affected by changes in a (unless c is set equal to a , in which case word-cue sampling is affected because c changes; see below). However, recovery probabilities depend directly on a . Thus a strongly influences both the overall level of recall and the dependence of recall on presentation time.

The value of a should depend on the coding or rehearsal technique used, on the nature of the words, and on the distinctiveness of the list context.

Parameter b. The interword strength per second of joint rehearsal time was estimated to be .10. The value of this parameter directly influences recovery probabilities whenever a word is used as a cue and an image is sampled that had been rehearsed with the cue word. However, if a word cue is used, b also affects the sampling probabilities. The reason for this depends on the three types of samples: It is possible to self-sample, to sample a word that had not been rehearsed with the cue word, or to sample a word that had been rehearsed with the cue word. As the value of b rises, the probability of the last of these possibilities rises.

The value of b should depend on context distinctiveness to a small degree but should depend primarily on the ease with which the list words may be coded together. For example, increases in the similarity of the list words to each other should probably increase the value of b for that list.

Parameter c. The word's self-sampling strength per unit of rehearsal time was set equal to the context-to-word strength, a . The self-sampling strength gives the relative tendency for a word to sample its own image, and it seems reasonable that this tendency

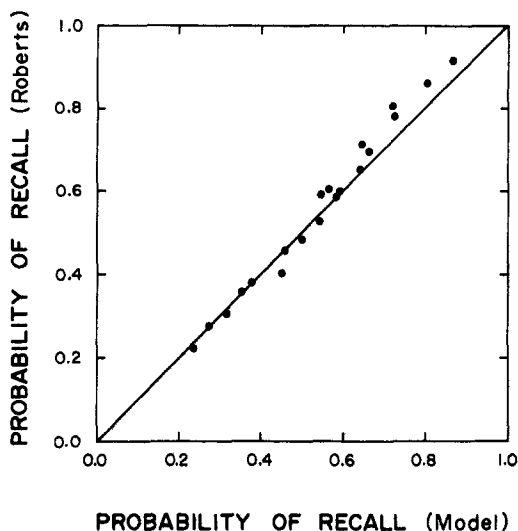


Figure 6. Comparison of the predicted and observed probabilities of recall, corresponding to each of the 20 points in Figures 4 and 5.

should grow with rehearsal time. Setting c equal to a accomplishes this goal, but there is no real reason to believe these parameters should be equal (rather than, say, linearly related). If c is raised for fixed a , then self-sampling rises and recall performance drops (though recognition performance would rise). For the recall tasks to which the model has been fit, the value of c does not seem to be crucial. Evidence about the relationship of c to a can best be obtained from recognition tasks, but for the present no harm is done by setting c equal to a .

In general, we expect c and a to vary in similar ways as the experimental conditions are changed.

Parameter d . The residual interword strength value for words not rehearsed together was set equal to .02. The value was chosen on the basis of applications of the model to data from paired-associate paradigms (not discussed in this article). This value is low relative to a , b , and c , as should be the case. In fact, the precise value is not important for the present application: Almost identical predictions of Roberts' data are obtained when d is set equal to .0, the value of a is raised to .12, and the other parameters are unchanged.

The effect on performance of raising d is complex. Sampling will be directed away from words rehearsed with a cue word, which could hurt performance in a cued-recall or recognition task. However, in the present free-recall task, performance is improved slightly because more retrieval routes are made likely, because self-sampling is reduced, and because recovery probability is raised.

It is presumed that d reflects preexperimental factors and should be higher for words that are already strongly related, especially if that relationship was formed in a context similar to the test context. Also, factors that affect the value of b should usually affect d similarly.

Parameters e , f , and g . The increments in strength following recall were all set equal and then estimated, attaining a value of .70. The particular value chosen is not very important for the present application. However, an increase in these incrementing values does lower overall recall levels, so they must be estimated jointly with a and b . Incrementing is included in this application and set to a fairly high value to achieve consistency with applications to other paradigms, in which incrementing is essential to predict the data (see Section I, *Discussion*).

The amount of incrementing is supposed to reflect the amount of attention given to any word as it is recalled and after it has been recalled.

Predictions for Free Recall

The predictions of SAMS with these parameter values are shown in Figures 4 and 5 in the lower panels. Since direct comparison of observed and predicted points is difficult in these figures, Figure 6 shows a comparison of the predicted and observed points for all 20 conditions. Clearly the fit of the model is quite adequate.

Discussion

Our parameter estimation procedures make it clear that this model is far more powerful than it needs to be to fit free-recall data. Many of the model's processes and

parameters are included to deal with data from other paradigms that we do not have the space to discuss here. However, we will present a brief survey of certain findings that help illuminate the roles played by these processes.

First, consider the search stopping criterion. One problem with a search model such as that of Shiffrin (1970) was its dependence on the number of samples made from memory: The number of samples was equated for all lists. Perhaps such an assumption is defensible if the recall period is short and equal for all lists. However, in almost all studies the recall period is set long enough that the subject ceases attempts to recall before time expires. The present model incorporates a stopping criterion based on the progress of retrieval during any given search. Furthermore, the criterion is a subject choice that should be alterable by the experimenter. Roediger (1978) and Roediger and Thorpe (1978), for example, have examined both the time course of retrieval and the effect of instructing subjects to continue recall attempts beyond the point at which retrieval normally ceases. In both cases, our model predicts the data with no changes in assumptions and only minor changes in parameter values. It is worth noting that our assumption of $K_{MAX} = 30$ causes recall to cease during a period when new recalls are occurring at a very slow rate relative to the rate occurring earlier in the search. This comment is supported by the observation that an alternative stopping rule, in which search ceases when 12 consecutive failures occur, produces predictions almost identical to those obtained for a rule of 30 total failures.

Consider next the purpose of rechecking and the effects of this process on the predictions. Rechecking was included initially so that it could be argued that the subject had exhausted all available retrieval routes. Rechecking is not essential to predict Roberts' (1972) data; the elimination of rechecking can be handled without difficulty by increasing the value of K_{MAX} or raising the values of a and b . Rechecking does play an important role in other situations, however. For example, without rechecking, cumulative recall functions can be predicted to grow

to asymptote too sharply; rechecking provides new retrieval routes that allow total recall to continue to grow (albeit slowly) for long periods of time (Buschke, 1975, Roediger & Thorpe, 1978).

Consider incrementing next. Although not really crucial for predicting Roberts' data, this process is essential for predicting a number of effects from other paradigms. For example, two consecutive tests of the same list may be given, one by the technique of free recall and one by some sort of cued recall (e.g., category cues or part-list cues), in either order. Quite large order effects have been found in our own research, and these can be explained quite well by the effects of incrementing. A somewhat smaller effect is that of test order of categories in cued recall of categorized word lists. Slight drops in recall are found for categories that are delayed in testing (e.g., Roediger, 1973; Smith, 1971). This effect also is well predicted by incrementing. The reason why incrementing explains these effects is straightforward: Recall leads to incrementing that increases the sampling probability for the items already recalled.

Interword associations and their use in retrieval are a cornerstone of SAMS, but even these are not needed to predict Roberts' data. Such associative routes are needed, however, to predict many results, including those concerning paired-associate paradigms and the effects of variations in interword association values. Perhaps most intriguing is the manner in which such interword retrieval routes explain the part-list cuing effect, which is discussed in the next two sections.

Finally, consider the residual strengths between list words not rehearsed together. We include such retrieval routes in part because it seems logical that they should be present. However, many results show the need for such routes. For example, if a list of word pairs is studied, with instructions to form codes for each pair, and one member of a pair is presented at test, recall probability depends strongly on list length (Raa-iymakers & Shiffrin, 1980). This effect is due within the model to the assumption that the cue word is associated not just to the

response word but also, residually, to each word on the list.

We realize, of course, that no serious support for our present model can be obtained from our fit to Roberts' free-recall data. However, the applications to part-list cuing discussed in the next two sections help show the value of the present approach.

II. Part-List Cuing

A Survey of the Literature

One of the more intriguing results obtained in the free-recall paradigm is the so-called "part-list cuing" effect, first reported by Slamecka (1968). The effect is important because it is counterintuitive and seemingly inconsistent with many theories of human memory.

In the prototypical part-list cuing experiment, subjects are given a list of words for study and are later tested for recall of these items in one of two ways: One group of subjects, the control group, simply recalls as many items as possible; that is, this group is given the usual free-recall instructions. A second group is given a randomly chosen subset of the list items as cues, supposedly to aid retrieval of the remaining items. The recall by this latter, part-list cued group of the remaining or *critical items* (also called *target items*) is then compared with the performance of the control group on the same items. This paradigm was devised by Slamecka (1968) as a rather direct test of the role of interitem associations in recall. Slamecka reasoned that any theory that assumes that interitem associations are used in recall should predict that at least some of these cues would facilitate recall of items that would otherwise not have been recalled. Failure to find recall facilitation should (according to Slamecka) lead to the conclusion that interitem associations are not formed during study or not used during recall.

The result that is obtained quite reliably in this paradigm and in a number of variants shows that recall of the critical items is not better for the cued group than for the control group (even when short-term memory effects are controlled). In fact, there is usually a slight negative effect, with the control

group showing superior recall (Roediger, Stollon, & Tulving, 1977; Slamecka, 1969). What is surprising about this result is not so much the small superiority of the control group, but the fact that no reliable cuing facilitation is found.

Most conventional theories assume that during study of a list of words a network of interitem associations is formed. A model postulating that the list cues simply give a number of additional entry points into the network should predict a large advantage for the cued group. For example, Anderson's FRAN model, which makes exactly this assumption, predicts much better critical word recall by the cued group than by the control group (Anderson, 1972, p. 362).

In the early experiments on part-list cuing (Slamecka, 1968), a rather large advantage was found for the control group. However, Slamecka (1968, Experiments V and VI) showed that this was because the control group had the advantage of being able to report the items still in short-term store, whereas the cued group had to process the list cues. When this artifact was eliminated, no significant differences were observed between the two conditions. Moreover, the lack of a positive cuing effect did not seem to depend on the interitem associative strength. For example, in one of Slamecka's experiments (Slamecka, 1968, Experiment VI), three types of lists were used: a list of 30 rare words, a list of 30 common words, and a list composed of the word *butterfly* and 29 of its most frequently named associates. The mean number of target items recalled by the control group was 5.58 for the rare word list, 7.04 for the common word list, and 8.50 for the butterfly list. For the cued group these means were 4.70, 6.79, and 8.97, respectively. The difference between the cued and the noncued conditions was not significant ($F < 1$), nor was the interaction between this cuing factor and the type of list ($F < 1$). The differences between the three lists were, however, highly significant.

Thus, the part-list cuing effect does not seem to depend very much on the interitem strength. Although the control group advantage does seem to diminish somewhat or even reverse slightly with higher interitem strength, no large positive effect of cuing

appears, even when the variation in interitem strengths between conditions is enough to make total recall differ by a factor of nearly two. Such a result has been thought by previous investigators to be at odds with the typical associative network theories.

Similar results have been reported by a number of other researchers. Roediger et al. (1977) reported a control group advantage that increased slightly with the number of cues provided. They also observed that the advantage of the noncued over the cued condition did not disappear or reverse even when very long recall periods were used (up to 10 minutes).

Slight positive effects of cuing (often not significant) have been reported, but only in special conditions. Anderson (1972) and Wood (1969) reported a very slight positive effect after multitrial free-recall learning. In a similar experiment, Slamecka (1969) observed a slight negative effect of cuing. Allen (1969) presented subjects with a list composed of contiguous pairs of items, which were judged to be either related or unrelated. After an initial recall for both groups, he cued the subjects in one group with one member of each word pair and found greater recall for the cue group, especially with related pairs. Note, however, that this procedure is quite different from the normal procedure, in which a random subset of words is given as cues. The only paradigms in which a reasonably large advantage appears for the cued condition are retroactive interference paradigms. Blake and Okada (1973) reported an experiment in which subjects were given 10 study-test trials on a 6-item list, followed by 10 study-test trials on a second, interpolated list. Following this interpolated learning they were tested on the first list. In this experiment, cuing on the final test trial did have a positive effect. Similar results were reported by Basden (1973).

Roediger (1974) has summarized a number of experiments with categorized lists in which the number of cues per category varied. In these paradigms, cuing may help in retrieving categories when there are too many categories in the list to be retrieved unaided, but cuing has no beneficial effect at all on within-category recall, that is, on the number of words recalled per category

of which at least one member was recalled. In fact, a slight negative effect is consistently found in within-category recall. Results consistent with this generalization have been reported by Hudson and Austin (1970), Mueller and Watkins (1977), Roediger (1973, 1974), Rundus (1973), Slamecka (1972), and Watkins (1975).

Moreover, it is usually found that the negative effect of cuing on within-category recall increases slightly with the number of cues provided (Mueller & Watkins, 1977; Roediger, 1973, 1974; Watkins, 1975). Roediger (1974, Table 1) summarized a number of experiments that show that the higher the proportion of items from a particular category that are given as recall cues, the lower the probability of recalling the remaining items from that category. For example, Rundus (1973) obtained recall probabilities for the remaining category members of .36, .34, .29, and .28 for cuing with 1, 2, 3, and 4 category instances, respectively.

Explanations for the Part-List Cuing Effect

The basic part-list cuing result—the lack of an advantage for the cued group—is interpreted by researchers in this area as evidence against the role of interitem associations in free recall. Slamecka (1972), Rundus (1973), Roediger (1974) and Watkins (1975) assume only vertical or hierarchical associations. Thus, Slamecka (1972) writes:

A truly associative theory of learning, one which identifies learning with the active elaboration and strengthening of interitem linkages, cannot, in our opinion, convincingly incorporate these findings without abandoning a central feature of the term "association." The unique force of the proposition that A and B are associated is that presence of A has the power to elicit B; that A can "pull out" B. (p. 331)

This conclusion is shared by most other researchers in the area:

It should be noted that the results . . . contradict predictions of theories postulating direct associations between items. If items were directly interlinked in memory, it should be the case that presentation of some items as cues would increase the likelihood of other items being recalled. (Roediger, 1974, p. 265)

However, we will show in the next section that this reasoning is quite incorrect. Despite Slamecka's contention that the part-list cuing

effect shows that interitem associations are not used in recall, the effect is completely consistent with a model that relies heavily on the use of interitem retrieval routes. In fact, we shall see that the effect is virtually an inevitable consequence of models like SAM.

First, let us consider the explanations that have been developed by theorists who have decided to abandon the role of interitem associations. These explanations have tended to postulate hierarchical, or "vertical," associations in place of interitem, or "horizontal," associations. Such theories have not been worked out in detail, so it is not clear how they propose to handle the part-list cuing effect. One hypothesis would assume that a retrieval path from an item to a node at another level and then back to another item at the initial level is either difficult or impossible to follow: Thus a word cue at one level will not effectively cue another word at the same level. Whatever the merits of this hypothesis for part-list cuing, it would certainly have difficulties handling many standard results. These include the effects of input order on output order (Anderson, 1972; Kintsch, 1970; Shiffrin, 1970) and the effects of associability of items presented contiguously in the list on the level of recall (Glanzer & Schwartz, 1971). In fact, it is difficult to see how this hypothesis would handle single-trial paired-associate recall: Why should a stimulus cue increase the probability of recall of the response member unless there exists a useful retrieval route between the two?

Although all of the theories proposed to explain the part-list cuing effect make similar assumptions in order to explain the lack of a positive effect, they do differ somewhat in whether they can explain, or in the way they explain, the slight negative effect of cuing and the dependence of the cuing effect on the number of cues presented.

Slamecka (1972) proposed a generation-recognition account of categorized recall: He assumed that there is no strengthening at all during study of the category-name-to-instance associations. Therefore, his theory does not predict an effect of the number of cues provided to the subject. In fact, it does

not predict a negative effect at all. Although this accords with some of Slamecka's data (Slamecka, 1968, 1972), Roediger (1974) has shown that in the majority of studies there is an increasing negative effect as the number of cues rises. Two related theories have been proposed to deal with this phenomenon.

First, Rundus (1973) and Roediger (1973, 1974) have suggested that presentation of the list cues involves an implicit retrieval of these items. It is assumed that this retrieval results in a higher associative strength of each retrieved item to the retrieval cue (the category name). It is further assumed that retrieval involves a sampling-with-replacement process, so that this strengthening of the list cues leads to a lowered probability of sampling critical items. The model also includes the assumption that in noncued recall, in which the subject first has to search for the category names to be used as retrieval cues, the probability of sampling a particular category name is proportional to the sum of the associative strengths of the category exemplars presented in the list to that category name. Data consistent with this latter assumption have been presented by Parker and Warren (1974) and Roediger (1978). This model could presumably also explain the positive effects of cuing reported by Blake and Okada (1973) and Basden (1973) by assuming that in multitrial free-recall learning a subjective hierarchical organization is formed and that retroactive interference leads to a decrease in the probability of generating the subjective retrieval cues.

Second, Watkins (1975; see also Mueller & Watkins, 1977; Watkins & Watkins, 1976) has proposed a cue-overload interpretation of part-list cuing. According to this explanation, recall is mediated by retrieval cues that are subject to overload: The probability of recalling any particular item decreases with the number of instances associated to the retrieval cue. Thus, in this case the part-list cuing effect is attributed to the usual (but in this model, unexplained) list-length effect. Watkins (1975) assumes that re-presentation of the list-cues at the time of testing is functionally equivalent to presentation of a new category instance. This

hypothesis is supported by data of Watkins (1975) and Mueller and Watkins (1977) showing that recall of the remaining category instances is decreased not only by presentation of list cues but also by presentation of extralist cues, category instances that were not on the original list. Unrelated cues, however, have no effect on recall (Mueller & Watkins, 1977, Experiment I). Note that this explanation is quite similar to the one proposed by Roediger (1973, 1974) and Rundus (1973), if the latter model were modified to incorporate the negative effect of extralist cues.

Within the context of search theories of either the type represented by SAMS, the type described in Shiffrin (1970), or the general class discussed in Section IV, several explanations should be considered at this point. Since recall depends on the length of the search, it is possible to predict the part-list cuing effect simply by assuming that the subject searches longer in the control condition than the cue condition. One possibility is to allow the stopping criterion (K_{MAX} in SAMS) to vary between the two conditions. Another possibility is to count failures in such a way that they accumulate faster in the cue condition. For example, each sample and recovery of an image from the list of provided cues could be counted as a search failure, leading to earlier search cessation in the cue condition. Such mechanisms may possibly contribute to the observed effect, but certain evidence makes it seem unlikely that they are operating. The main problem is that cumulative recall as a function of time shows a consistent advantage for the control condition, even for times when all subjects in both conditions are still searching memory. This is shown most persuasively by Roediger et al. (1977), who instructed subjects to continue attempting memory search for a period of 10 minutes (by which time very few new items were being retrieved by either group). On the whole, it seems doubtful that different stopping criteria provide a general explanation of the part-list cuing effect. Such possibilities are explored further within the context of the SAMS model in the next section.

Finally, there is one explanation that ac-

cepts the formation of interitem associations during study. This is the editing/disruption hypothesis proposed by Basden, Basden, and Galloway (1977). They assume that

the part-list cues disrupt any intracategory organization that may have taken place during acquisition. Since the order in which the items are presented as cues is very unlikely to coincide with the order in which the subject would recall the category members, part-list cuing may constitute an interference paradigm. (p. 104)

Note that this hypothesis is still similar to the other explanations, since it is assumed that the interitem associations cannot be effectively used in recall. This particular explanation, however, seems to run into difficulties when applied to the positive cuing effects observed by Tulving and Pearlstone (1966), who presented category names as cues in a random order, and Tulving and Osler (1968), who presented extralist cues for noncategorized lists. Moreover, no reasons are given why the items should only be recallable in one particular order.

In summary then, the extant explanations have certain drawbacks, and most of the major explanations agree with Slamecka's (1968, 1969) original conclusion that the lack of a positive effect of cuing proves that interitem or horizontal associations play no important role in free recall. This of course presents us with a theoretical puzzle: What possible reason could there be for horizontal associations not being stored? In the next section we will present a model and an explanation for the part-list cuing paradigm that does utilize horizontal associations and thereby eliminates this puzzle.

III. Application of SAMS to Part-List Cuing

First, SAM will be applied to part-list cuing, with the interitem strength, b (i.e., interword similarity), as a variable. Then a number of factors that do not affect the basic predictions will be mentioned and illustrated. Next we shall endeavor to explain the factors that do cause our model to predict the effect. We shall then apply the model to several additional findings from the literature, namely, the effects of number of list cues, the effect of categorized lists, and the

effect of various possible types of list cues. Finally, our explanations will be compared to others in the literature.

Assumptions of the Model

The model proposed for part-list cuing is the SAMS model already fitted to Roberts' (1972) data, with a few slight changes to enable it to deal with the cued group. An idealized paradigm is used as a basis for the theoretical development in which N words are presented for study and followed by arithmetic to clear STS. The control condition is normal free recall. In the cue condition, M words randomly chosen from the list are presented to the subjects at the start of recall. The subjects are told to utilize the M words as cues to help them recall as many of the remaining ($N-M$) words as possible.

The model for the control condition is the SAMS model that was fitted to Roberts' study and has the same parameter values except that STS recall is not allowed, since the arithmetic task clears STS.

The assumptions for the cue condition are identical except for a few special assumptions that apply at the start of recall to take the list cues into account. We are of course careful not to make any assumptions that will introduce an advantage or disadvantage for the cued group for trivial reasons (such as an assumption that the cued group used a smaller value of K_{MAX} to determine when to cease searching). Also, we are careful to treat the list cues in the cue condition, which are provided by the experimenter, in the same way as the self-generated cues in the control condition.

In particular, it is assumed that the subjects act in accord with the instructions and utilize the list cues during the search. At the start of recall, each list cue is used in turn as a retrieval cue along with context. Each such cue set is used until L_{MAX} failures are reached, and then the next list cue is used, and so forth. During this phase, the subject does not use any recalled words as cues, but instead saves them for later use. In order not to introduce a disadvantage for the cue group, the first successful recovery of the image of any list cue, whether due to self-

sampling, context sampling, or word-plus-context sampling, is not counted as a failure (see Section II, *Explanations for the Part-List Cuing Effect*.) Only recoveries of the image of a list cue after the first are counted as failures. This assumption is consistent with the retrieval assumptions for the control group, since any word's first recovery is a success for that group. The usual incrementing and conditionalization assumptions apply during this phase of retrieval for the cue group.

When the list cues have each been used for L_{MAX} failures, and if K_{MAX} has not yet been reached, then the subject begins normal search, just as at the start of recall in the control condition. This continues until a criterion of K_{MAX} total failures of any kind is reached, at which time Phase 1 of the search ends for both groups.

Following the K_{MAX} stopping point, both groups carry out a final rechecking in which all successfully recovered words are used as cues. Thus, recovered list cues are utilized during this phase, but unrecovered list cues are not. (This last assumption ensures that the cue group receives no advantage due to an artificially extended rechecking period. Alternative rechecking assumptions will be discussed later.) During rechecking, the usual assumptions apply, just as in the model described earlier.

Predictions of SAMS for the Part-List Cuing Effect

We now turn to the predictions of the model. We assume that using a highly similar list of words (as in Slamecka's "butterfly" list) is equivalent in our model to using a high value of the interitem parameter, b . Presumably, semantically similar items are easier to encode together. Therefore we derived predictions for six different conditions that differed only in the values assigned to b . These values ranged from very low to very high, as shown in Figure 7. Presentation time was set to 2 sec per item, list length was set to 30, the number of cues in the cue condition was set to 15, and all other parameters were those used in the fit to Roberts' data (the residual strength, d , was set in each

condition to one fifth of the value of the interitem strength, b ; this was the ratio for the b and d values in the fit to Roberts' data). NSIM was set equal to 1000.

Figure 7 gives mean critical words recalled for the cuing and control conditions for each of the six values of interitem strength. The predictions are interesting in two respects. First, the control condition has a small but systematic advantage over the cuing condition. Second, the size of the control group advantage is fairly consistent over the range of interitem strength values, decreasing slightly with higher values of b . Note that total recall rises from about four to eight items over this range of values of b without much altering the part-list cuing effect. Of course, this pattern of results is very similar to that found by Slamecka (1968; and see the discussion in Section II. of this article).

Note especially that extensive use is made of the associative pathways between words, which is indicated by the fact that predicted recall almost doubles over the range of interitem strengths in the figure. Certainly, then, the part-list cuing effect cannot be used to argue that interitem associations are either not stored or not used in retrieval.

An important fact about the model that needs to be emphasized and kept in mind throughout the following sections is that the control group advantage in Figure 7 occurs in spite of a fairly large factor that aids the cue group. The model incorporates a recovery rule which insures that the probability of recovery following sampling of an image with a word-plus-context cue set will be higher than when that same image is sampled with a cue set consisting of context only. Since the cue condition utilizes relatively more word-plus-context cue sets and less context-only cue sets, it achieves a considerable advantage. This is seen most clearly when the recovery probabilities for word-plus-context and context-only sets are equated. As a demonstration, we replaced the recovery equations in SAMS with a single probability of recovery, .75, which applied regardless of the context set used or the strengths involved (the usual condition-alization rules about successive sampling of




Figure 7. Predictions of the part-list cuing effect for an idealized paradigm as the interitem strength parameter, b , is varied. (List length is 30; number of cues is 15; presentation time per word is 2 sec; d is set to equal .2 b . Other parameters are as in Figure 4. These values hold for the following figures, unless indicated differently. For each strength value, the control condition is slightly superior, a pattern of results similar to those of Slamecka [1968].)

the same image still applied). All other features of the model remained as in Figure 7. When this was done, the control group advantage rose by about .5 items.

The important implication of this finding is that all the predictions we shall be discussing in the following sections are occurring in the presence of a recovery factor aiding the cue group. Thus even predictions of equality of the control and cue conditions, which will occur in several variations below, are actually indicating the operation of substantial factors favoring the control condition (since for equality to occur, the cue condition recovery advantage must be overcome).

Factors That Do Not Explain Part-List Cuing

Ineffectiveness of word-plus-context searching. Consider first whether searching with word cues plus context cues is superior to searching with context only. The answer may be seen most clearly when rechecking is eliminated from the model. The predictions for such a case are presented in Figure




Figure 8. Predicted words recalled in the control condition, without rechecking, for normal search as a function of interitem strength (solid line) and for search with context cues only (dashed line).

8. For both conditions normal free recall is assumed (no list cues are used). The context-plus-word cue curve uses Roberts' parameters and shows how total recall depends on the value of the interitem strength, b . The context-only condition assumes that all samples are made with context only. The results show that the possibility of word-plus-context sampling causes recall to be superior at high values of b , but the reverse is true at low values of b . This is to be expected, since low values of b will result in word cues sampling themselves (since $c = a$) rather than other list items. The two conditions are about equal for b near the value estimated for Roberts' data. However, this factor evidently has nothing to do with the part-list cuing effect, since the part-list cuing effect is stable and almost equal over the same range of b values that produces the large differences in Figure 8 (compare with Figure 7).

Incrementing. The process we have termed *incrementing* has been proposed as an explanation of the part-list cuing effect. It is possible that the list cues become strongly associated to the search cues, particularly context cues, during the initial phase of the search. Then the later search tends to sample list cues to the detriment of critical items. Explanations of this general sort have been proposed by Rundus (1973), by Roediger (1973, 1974; see also Roediger et al., 1977) and, albeit differently phrased, by Watkins (1975). Since incrementing is

built into our model, could it be accounting for the effect? Figure 9 shows a set of predictions just like those in Figure 7, except all the incrementing parameters have been set to zero. Clearly, overall recall is increased when incrementing is eliminated, but the advantage of the control condition is unchanged. Thus the SAMS incrementing process and the part-list cuing effect are virtually independent within this model. The reason is that incrementing has equally deleterious effects on the cue and control groups. (The effect of incrementing the list cues before retrieval begins will be discussed later.)

It should be noted that Crowder (1976) mentions that incrementing is not needed within a search model to explain the part-list cuing effect. This prediction was derived in the context of a model assuming that all retrievals of list cues (including the first) are counted as failures—an assumption we do not make. Without this assumption, his line of reasoning is invalidated.

The choice of stopping rule. Consider variations in our assumptions about the stopping rule. In the first variation, we changed the criterion from " K_{MAX} total failures" to " K_{MAX} consecutive failures." Figure 10 is comparable to Figure 7, except that the consecutive failure rule is used, with $K_{MAX} = 10$. Clearly, the choice of stop rule has little to do with the part-list cuing effect. Next, for both types of stopping rules, various values of K_{MAX} were tried, since it might be argued that SAMS' predictions only hold true when




Figure 9. Recall as a function of interitem strength. (Predictions use the same model as Figure 7, but all incrementing is deleted [$e = f = g = .0$].)




Figure 10. Recall as a function of interitem strength. (Predictions use the same model as Figure 7, except that the total failure stopping rule for Phase 1 of the search has been changed to a consecutive failure stopping rule, with $K_{MAX} = 10$ consecutive failures.)

searching ceases quickly (i.e., too soon). Figure 11 shows the predictions for various values of K_{MAX} for the total failure rule, with b set to .10. Clearly the effect and its magnitude are independent of the stopping criterion. (Similar results obtain for the consecutive failure rule.) Note that for $L_{MAX} = 3$ and $K_{MAX} > 45$, all the list cues are used in Phase 1 of the search. The fact that some are skipped for $K_{MAX} = 30$ makes very little difference.

Despite the predictions in Figure 11, one might still be concerned that stopping criteria are responsible in some way for the effect. Thus in Figure 12 we deleted the stopping rule, and simply graphed cumulative recall for both groups as a function of the total number of samples made. (It is assumed that rechecking takes place every 100 samples; without rechecking, very little additional recall for either condition would take place late in the search, since most images would have been sampled to the context cue at least once.) It is clear that recall increases at a rapidly decreasing rate, even with rechecking. The control group retains its advantage throughout, although the difference decreases. Roediger et al. (1977) obtained results somewhat more like Figure 11 than Figure 12, which perhaps suggests that their subjects employed a stop rule with a large criterion, but it must be admitted

that both figures would fit their pattern of results quite well.

The choice of probe cues. It may be noted that the search assumptions for the two conditions differ slightly, since the control group is assumed to switch cues at once whenever a new word is recalled, whereas the cue group uses each experimenter-provided cue until L_{MAX} failures accumulate. In fact, if each time a word is recalled, the cue group is made to move on at once to the next list cue, the two conditions become virtually equal. Conversely, if the control group is forced to stay with each word cue until L_{MAX} failures accumulate, saving all recalled words for later use, whenever K_{MAX} is reached, then again the two conditions become almost equal. In neither case, however, is there a cue group advantage.

Assumptions differing for the two groups. Several assumptions can be made that will favor one condition or another for the rather obvious reason that the two groups are treated differently. For example, assuming that all recoveries of list cues are failures causes the cued condition to reach K_{MAX} more quickly than the control group. The advantage of the control group naturally increases with this assumption (by about .15 words). Consider another example: If all list cues, whether recovered or not, are used in rechecking, then the cued group spends more

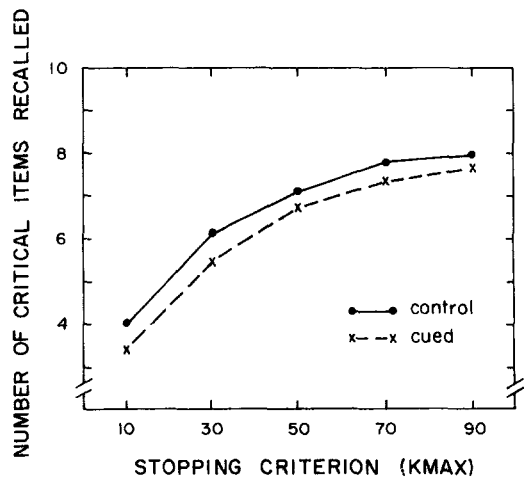


Figure 11. Recall as a function of stopping criterion. (Predictions use the same model as Figure 7, but with interitem strength $b = .10$.)




Figure 12. Cumulative critical words recalled as a function of total number of samples. (Predictions use the same model as Figure 7, but with interitem strength $b = .10$ and stopping rule deleted. Rechecking occurs every 100 samples.)

time rechecking than the control group and naturally gets an advantage from this fact (the cued group now gets an advantage of about .5 words). Even without this assumption, rechecking tends to favor the cued group very slightly. When rechecking is eliminated altogether, the control group advantage increases by about .40 words.

The lesson from the manipulations that have been tried seems clear: If we remove every possible simple factor that might favor the control group, apart from the basic structure of SAMS itself, but retain factors that favor the cue group, such as a recovery advantage when a word cue is added to context, and a new chance to recover whenever a new cue is used to sample a previously unrecovered image, at best we bring the cued group up to the level of the control group. There are clearly some fundamental, inherent characteristics of the SAMS model that produce the part-list cuing effect. These are considered next.

An Explanation of the Part-List Cuing Effect

By now, the reader must be wondering which factors are responsible for SAMS' prediction of the part-list cuing effect.

Interword cuing in the control and cue conditions. First and foremost, much of the

mystery is removed when it is realized that both the control and cue conditions involve extensive cuing by words. Self-generated cues are utilized in the control condition, whereas experimenter-provided cues are utilized first in the cue condition. The extensive interitem associative cuing in both conditions tends to make overall performance in the two conditions quite comparable.

This reasoning suggests that when the ability to generate cues, using context, is poor, the control group will be inferior. Figure 13 shows the predictions as a , the item-to-context strength, varies. At very low values, there is a considerable advantage for the cued group (since the control group can seldom recall anything using context alone), but as a increases in value, the usual part-list effect quickly appears and even grows somewhat. One might expect, then, that a cuing advantage could be experimentally produced simply by choosing a task in which context-to-item strengths are low at test, that is, a task in which normal free recall would give very low performance levels. Such reasoning could well explain the positive effects of cuing observed by Blake and Okada (1973) and Basden (1973) in a retroactive interference paradigm.

The selection of list cues. In order for the part-list cuing effect to be seen, it is extremely important that the subject be given a random sample of items from the list as cues. Otherwise, the cues may turn out to




Figure 13. Recall as a function of item-to-context strength parameter (a). (Predictions use the same model as Figure 7, but interitem strength $b = .10$ and $c = a$.)

be extremely useful. It is true in the model and has been found empirically that cues considerably increase the recall of those words that are strongly associated to the cues—it is the other words whose retrieval is harmed (Roediger, 1978). This is nicely illustrated in a study by Twohig (Note 1). Twohig presented subjects with a list of categories that were each composed of four pairs of associated words. Subjects in the cued group were given two words (not from the same associated pair) from each category as list cues. There was no positive effect of cuing, averaging over all critical items in a category. However, the conclusion that the interitem associations were ineffective would be quite inappropriate, since the cues did lead to a substantially higher probability of recall for the items paired with the cues. This positive cuing effect was, however, counterbalanced by a negative effect on the recall of the other category members.

In fact, this reasoning suggests that cued recall could be quite beneficial if the experimenter could arrange to provide one good cue from each of the subjective "groupings" that the subject has stored in memory. This is easiest to demonstrate when a categorical structure is built into the list. In this case, cuing with one word from each category is very beneficial (due to more categories' being accessed; see Slamecka, 1972; Tulving & Pearlstone, 1966). This factor can also help explain Allen's (1969) finding that cuing was helpful when the cues consisted of one member of each of a number of word pairs judged to be related. However, if the word cues are selected randomly (i.e., without regard to the categorical structure), then there is no increase in recall (Kintsch & Kalk, Note 2). Such results are a problem for those theorists who use the part-list cuing effect as evidence against the role of interitem associations in recall. Similar reasoning applied to these results should lead them to the conclusion that in such experiments hierarchical or vertical associations were not used, an assumption that probably none of them would like to embrace.

Clearly, then, appropriate selection of cues can lead to an advantage or disadvantage, depending on the basis for selection. In the studies under consideration, however,




Figure 14. A simplified associative network for a 12-word list stored as four triads. (The six experimenter-provided cue words have images denoted by Q. The six remaining critical items are denoted by I. The arrows denote associations between images. Each image has an association to context, which is not depicted. A context sample can access a triad rich in critical items [e.g., the triple-I triad]. The cue-word-plus-context samples can only sample triads relatively impoverished in critical items, since each such triad must contain at least one cue.)

random cue selection was used. Given that this is the case, an explanation for the cuing deficit needs to be established. We turn to this explanation next.

The effect of associative clustering. The strongest and most crucial factor favoring the control condition depends on the nature of sampling from a subjectively clustered associative network. It can be shown that the control group's sampled clusters will be relatively richer in critical items than the cued group's sampled clusters (most of which will contain at least one cue word).

The basic idea is illustrated in simplified form by the associative network depicted in Figure 14. It is assumed that the words are interassociated in groups of three (triads), with all other interword associations being negligible. It is further assumed that sampling of an image from a triad leads that member—and immediately thereafter, both other members of that triad—to be recalled. During a fixed period of time, assume that the control and cued groups sample an equal number of different triads (a simplification for the sake of the argument). The cued group's sampled triads will all contain a minimum of one cue word and hence a relatively small number of critical words. The




Figure 15. Recall as a function of strength value. (Predictions use the same model as Figure 7, but all entries in the test matrix at the start of retrieval are set equal to the same value, which is shown on the horizontal axis; thus r , a , b , c , and d are not used. Note the sizeable advantage for the cued condition.)

control group's sampled triads, on the other hand, will often contain no cue words and hence be relatively rich in critical words. Thus the retrieval structure forces the cued group to retrieve more list cues than the control group at the expense of retrieval of critical items, the measure on which the two groups are compared.

Keep in mind that for this simplified example, we assumed that both groups sample an equal number of different triads and hence recall an equal total number of words. If instead we assumed that an equal number of triads were sampled, including resamples, then it is clear that more different triads would be sampled and hence more total words would be recalled in the control condition. This factor would increase the size of the control group advantage even further.

In general, then, the cuing procedure leads to a sampling bias for the cued group. The cued group is more likely to sample list cues than critical items, whereas the control group is unbiased with respect to the two types of items. This sampling bias is so strong that it evidently can overcome other factors that aid the cue group, with the net effect that a higher number and proportion

of critical items is recalled by the control group.

Note that this explanation does not require that the subgroups in the structure be nonintersecting. A buffer process, for example, tends to produce high interword strengths in a region lying athwart the upper left to lower right diagonal of the strength matrix (Figure 1). This degree of grouping is sufficient for the present explanation to contribute substantially to the part-list cuing effect.

In addition to this sampling factor, there are a number of secondary factors that also contribute to the control condition advantage. One of these is a positive correlation that exists between a word's context strength and interword strengths. Such a correlation is produced by a buffer process but seems reasonable for any sensible storage system. (Note, however, that the SAMS buffer process does not lead to a particularly strong correlation.) That the correlation does contribute to the control group advantage is easy to demonstrate: in each simulation, after storage and before retrieval (for both conditions), the set of context-to-image strengths in the top row of the strength matrix of Figure 1 is randomly permuted. When this is done, the advantage of the control group is eliminated, and both conditions are about equal.

Note that both reasons for the superiority of the control condition depend on the development and use of a nonuniform interword retrieval structure. It is rather remarkable that the reason for the cue condition disadvantage in SAMS is the presence of the very associations that previous theorists have tried to rule out.

Additional evidence supporting our explanations of the part-list cuing prediction in the model may be obtained by eliminating the structure in the test matrix. The test matrix was homogenized by setting all strengths in the matrix equal to each other. In this case, both lines of reasoning described above were invalidated, and we expected a cue condition advantage due to its higher recovery probability. In fact a large advantage for the cue group appeared, as shown in Figure 15. Similar results, though not quite as extreme, were obtained simply by

raising the residual strength (represented by d) in the basic model. As the residual strengths approach the interword strengths, the grouping structure tends to be lost.

One final point should be emphasized. For the control group to do well, especially at high interword strengths (see Figure 8), it is necessary that an appropriate mixture of context-only sampling and word-plus-context sampling be utilized. Context-only sampling may tend to locate many clusters in memory but will not tend to retrieve efficiently the members of those clusters. Furthermore, the sampled words will tend to be more difficult to recover than when word cue strengths have been added to the context strengths. On the other hand, reducing context sampling to a minimum also reduces performance, since new clusters may not be found and inefficient word cues may be utilized over and over without success. Although the mixture of cue sets used in the model may not optimize performance for the control condition, it probably comes reasonably close. It goes without saying that the cue condition does not utilize an optimal sequence of cue sets. SAMS predicts that much better performance in the cue condition could be produced if searching went on for a period without considering the provided cues at all, and then the cues were utilized during a rechecking period (see Allen, 1969, and *An Application of SAMS to Delayed Cuing* below).

Applications of SAMS to Studies Varying the Number of Cues

To obtain predictions, we again imagined an idealized situation in which a 30-word list was presented and the number of words presented as cues was varied. It is easiest to compare the results when they are tabulated in terms of the probability of critical word recall (since total critical recall naturally depends on the number of critical words in the list). The objection might be raised that K_{MAX} can be reached before the provided cues are all utilized and that the number of such unchecked cues will grow with the number of provided cues. To eliminate this possibility, L_{MAX} was reduced to 2 and K_{MAX} raised to 50, so that all provided cues would




Figure 16. Probability of recall as a function of number of experimenter-provided cues. (Predictions use the same model as Figure 7, with interitem strength $b = .10$, $L_{MAX} = 2$, and $K_{MAX} = 50$.)

always be utilized in Phase 1 of the search. The other parameters were kept the same as in the fit to Roberts' data. Predictions are shown in Figure 16, which gives the probability of critical word recall as a function of the number of cues provided. Clearly, the model predicts a slight but almost linear decrease in recall as the number of cues increases. We have not explored the parameter space, but the magnitude of the decrease can presumably be controlled by choices of parameter values.

In the research literature, the effect of the number of provided cues is not entirely clear when uncategorized lists are used. Slamecka (1968) found little effect, but Roediger et al. (1977) did show an increasing deficit for the cued conditions as the number of cues increased. It is possible that subjects do not always use the provided cues, and firm instructions to read, study, and use the provided cues may produce this effect.

An Application of SAMS to Delayed Cuing

Allen (1969, Experiment 2) presented pairs of words—half related, half unrelated. After an initial period of 5 minutes of or-

dinary free recall, the cued group was given list cues, one from each pair. In the case of related pairs (high interword strength), the model predicts, and the data show, a substantial advantage for cuing. Allen also found a small cuing advantage for the unrelated pairs. To simulate this situation, we ran the uncued free-recall program until a criterion of $K_{\text{MAX}} = 45$. Then in the second phase, the cued group used each of the 15 list cues until a criterion of $L_{\text{MAX}} = 3$ was reached; the control group was simply continued free recall until an additional 45 failures accumulated. No rechecking was used. For Roberts' parameters, total predicted recall was higher for the cued group by .7 words, which was similar to Allen's findings for unrelated words.

Within SAMS, there are several reasons why delayed cuing should be helpful. For one thing, if there are any clusters of words not yet accessed, the list cues might provide entries to these clusters. Most important, however, are the new chances at recovery that the list cues provide. Late in any extended search, most images have been sampled by the context cue, and few new words are being recalled to serve as new word cues. On the other hand, many of the (delayed) list cues may not have been recovered yet. When these are used for sampling, they will provide new independent chances of recovery for images that were sampled earlier with other cues but not recovered. Allen's study provided evidence for this view. The list cues were separated into two classes: those that had been recalled in the immediate free-recall period and those that had not. The unrecalled cues were far more effective.

An Application of SAMS to Part-Category Cuing

Many of the studies of part-list cuing have utilized lists of categories of items and have varied the number of words presented as cues from each category. In addition, some of the studies have presented other types of cues. Applying SAMS to such paradigms involves essentially no new principles; in a sense each category is treated as a separate list, so most of the predictions of the preceding sections hold true.

A precise description of the application of SAM to the category case does involve a few additional assumptions, however, so these shall be described very briefly. It is assumed that each word image stored contains category information and each category label can be used as a probe cue (in much the same way, context is part of each image and can be used as a probe cue). In general, a given category cue will have strong strength values to words in that category and weak residual strengths to words in other categories.

In most studies the subject is told what category to recall at a given point in the search, so we assume that the probe set will consist of the context cue, the category cue, and possibly a word cue, if a word has been recalled from LTS or provided by the experimenter as a cue. Three kinds of experimenter-provided cues should be distinguished: *intralist cues*, words from the list in the to-be-recalled category; *extralist cues*, words not on the list but from the to-be-recalled category; and *extracategory cues*, words that may or may not have been on the list but are not from the to-be-recalled category. Watkins (1975) and Mueller and Watkins (1977) observed a negative effect of both intralist and extralist cues but no effect of extracategory cues.

The version of SAM for this case was developed and applied by Bruce Williams of Indiana University. It was assumed that the categories are blocked at input and that the usual buffer storage process operates, except that words from different categories are not rehearsed together (no interword strength is stored when words from different categories share the buffer). Instead, all words in different categories were given a common, small, residual strength. Words from the same category but not rehearsed together were given a (possibly different) small residual strength. Category-to-word strengths were a linear function of the time a word from that category stayed in the buffer, and category-to-word strengths for words in other categories were set to a small residual value. Finally, extralist cues were assumed to have a small residual strength to the category label and to list words in that category but negligible strength to words in other cate-

gories. All nonlist words were assigned a residual self-strength (which could be fairly high) and a residual context strength.

Retrieval of the list items from a given category was assumed to operate in essentially the same way as for a noncategorical list. Both context and category cues were used for every sample. If a noncategory word was sampled, it could be recovered but was not output, was not incremented, and was counted as a failure. In all other respects retrieval proceeded as for the noncategorized case, except there was no rechecking (when the model was applied with rechecking, no important changes in predictions resulted). For the cued conditions, the retrieval assumptions were again the same as before, also with the provisos of this paragraph.

Since it did not seem appropriate to fit the categorized case with the parameters for Roberts' data, we used the task and data from Watkins (1975, Experiment 1) and adjusted parameters roughly until a fit was obtained. The predictions and data for the control condition, the intralist cue condition, and the extralist cue condition for both two and four cues are shown in Figure 17. Clearly these predictions are quite adequate. One important misprediction was obtained, however. The model was applied to the extracategory cue condition of Mueller and Watkins (1977) and predicted a control condition advantage of about the same size as those shown in Figure 17 for the other cue conditions. The data showed no such deficit, even though the provided cues had been presented on the list in other categories.

We suggest that subjects given words from categories other than the category being tested will be very reluctant to use these as cues. Furthermore, we suggest that the predictions of SAMS would turn out to be correct, if only the subjects could somehow be induced to use these extracategory cues. A test of this possibility must await further research.

It is most interesting to note that despite the large number of parameters available in this category case, we could not find any combinations of values that would allow the prediction of the observed extracategory finding while simultaneously predicting the other findings (unless assumptions are

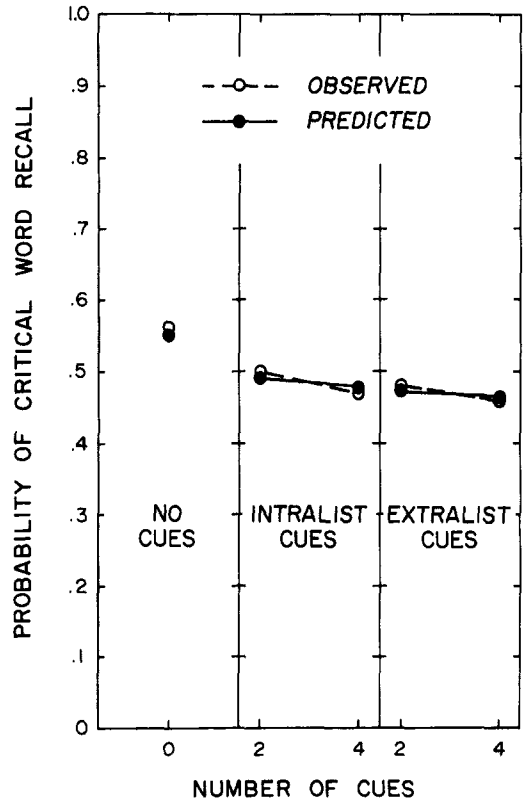


Figure 17. Predicted and observed probabilities of critical word recall in a category for the control condition and for the intralist and extralist cue conditions, each with two or four cues (data from Watkins, 1975). (The model is described in the text; parameters are list length = 36; words per category = 6; presentation time per word = 3 sec; $r = 4$; $K_{MAX}/category = 12$; $L_{MAX} = 3$; $a = c = .38$; $b = .38$; item-category strength per sec = .38; category-cue and list-word-cue residuals to words on list = .1; all increments = .36; residual strength of extralist cues to list items in same category = .03; and product of residual strengths when self-sampling an extralist cue = 2.2.)

changed between conditions). SAMS quite persistently produced a part-list cuing effect in the extracategory condition whenever one was predicted for the extralist condition. Thus we must not in this instance confuse number of parameters with testability. (Actually, the values of most of the parameters do not affect the qualitative predictions.)¹

¹ The explanations in Section III for the cuing deficit clearly do not apply in the case of extralist cues, since an extralist cue is not part of any stored cluster. The deficit in the extralist case results from two factors:

Conclusions and Comparisons With Other Models

SAMS predicts essentially all the phenomena of part-list cuing. It does so robustly under virtually all combinations of assumptions and parameter values and in fact cannot be made to predict the extracategory cuing result for this reason.

Most importantly, the basis for SAMS' prediction of the part-list cuing effect is the presence and utilization within a search theory of the very interword associative network that previous theorists have argued cannot be present or cannot be utilized. In short, it is proposed that extensive use of word cues is made in both the control and cue conditions and that the sampling superiority for the control condition outweighs the advantage during recovery of the cue condition. The control condition superiority is rooted in the nature of sampling from a network of (overlapping) clusters of words. When experimenter-presented cues are utilized, there is a tendency to gain access to clusters containing more critical items in the control condition than is the case in the cue condition.

The reader may wonder why we have expended so much effort in this section exploring the predictions of the model under alternative assumptions. It may seem that we are studying the workings of the model as much as the workings of the subject. There is some truth in this observation. In fact, we have attempted to deal with a problem that Smith (1978) has termed "the sufficiency/transparency trade-off." The problem is that as a model becomes more and more complex and gains power to fit the data, it becomes increasingly opaque to the external observers (even the model's creators). The general principles are lost in the forest of details that such theories come to contain. Our present model is less complex

than many; in fact, SAMS may be written in two pages of FORTRAN code. Even so, the basic principles underlying the predictions of the part-list cuing effect would be virtually impossible to ascertain, were we to present only the general model and the predictions. Our solution to the problem in this instance was an extensive search of the "assumption space" of the model.

The fact that SAMS predicts the part-list cuing effect does not invalidate other hypotheses that have been proposed. Even if SAMS is basically correct, other factors might be adding to or subtracting from the size of the effect. However, it should be kept in mind that if a SAM-like approach is adopted, these other hypotheses are not needed to handle the effect.

Let us consider briefly some of these alternative proposals. The possibility that incrementing of the context-to-cue strength takes place when each cue is first used is certainly worth considering (see Roediger, 1973, 1974; Rundus, 1973). We added this factor to SAMS to see what additional advantage for the control group would result. In particular, when each cue was first utilized in a sample, it was given an increment to context (the same increment that normally applies). Everything else in the model remained as before, and the predicted control group advantage increased by less than .05 words. Thus this factor contributes relatively little when added to SAMS (though it might be more powerful if the cue increment were higher than normal or in a different type of model).

There are many fairly trivial ways to produce a control group advantage, some of which are nevertheless plausible and worth considering even within the context of a SAM model. For example, the control group might search longer than the cue group. This would happen if every recovery of a cue, even the first, counted as a failure. Then the cue condition would reach K_{MAX} much sooner than the control condition. One difficulty with this explanation is the fact that it predicts a strong dependence on the number of cues. Additional cues do harm recall, but only to a slight degree (e.g., Roediger, 1974; Slamecka, 1968, 1972; Watkins, 1975).

In any event, it is not known what search

First, an extralist cue has a high probability of sampling its own image, since it is only weakly connected to list words. Second, because the extralist cues are weakly connected to list words, incrementing will have a disproportionately large effect, causing later searches to resample with high probability any word recalled earlier. These factors do not apply in the usual case of intralist cuing.

strategies subjects adopt; if subjects are induced to search longer (that is, continue despite more failures) in one condition than another, our model clearly predicts a corresponding performance change. If a large value of K_{MAX} is used for the cued condition, the amount of extra searching needed to significantly raise recall for the control condition can be quite large (since new items are hard to recall after a long time searching). Thus the differential in the K_{MAX} values might have to be quite large, perhaps too large to justify. Nevertheless, the possibility of differential stop rules in the different conditions cannot be ruled out. We comment merely that such hypotheses are unnecessary and inelegant.

A somewhat different explanation within the context of a search theory would postulate that self-generated word cues (in the control condition) are more effective than the experimenter-provided cues of the cue condition for the following reason: A retrieved image may contain word plus study-context information, whereas a provided cue may contain word plus test-context information. Therefore the retrieved word cues may have higher strengths to the other words in storage than the provided cues. Higher strengths could lead to higher probabilities of recovery for sampled images. Such a factor could indeed lead to a control group superiority. Whether this is an important factor is another question. Presumably, the retrieved cue advantage only occurs to the extent that the retrieved study context provides information not already in the general context cue. However, a subject might well accumulate study-context information as the search proceeds and incorporate it into the context cue. If so, the difference between a retrieved cue and a provided cue might be minimal, except possibly for the first few loops of the search. In conclusion, we admit that a difference in cue effectiveness due to context retrieval might be contributing to the part-list cuing effect. However, such an assumption is not needed for SAMS to predict the effect, as we have seen.

A proposed explanation of part-list cuing that has received considerable interest is the hypothesis that "vertical" or hierarchical, associations but not "horizontal," or interi-

tem, associations are stored in memory (e.g., Roediger, 1973, 1974; Rundus, 1973; Slamecka, 1972). We are not really sure how this hypothesis predicts the part-list cuing effect, unless retrieval paths between items on the same level of analysis are not allowed to traverse other levels of analysis, a hypothesis easy to rule out. In any event, there is nothing wrong with vertical or hierarchical associations, and they are already part of SAMS, as in the categorized recall situation. It is the absence of horizontal, interitem, associations that seems unusual. There seems little point in ruling out interword associations, since the success of SAMS shows this hypothesis to be unnecessary.

Consider, finally, the "cue-overload" hypothesis (Watkins, 1975), which proposes that the cues act more or less as additional list items or category items, thereby reducing performance for the same reason that words in longer lists are recalled more poorly (whatever that reason is). This effect presumably operates despite the temporal and contextual separation of the list cues from the list. In some ways this hypothesis is like the cue-incrementing hypothesis discussed above, since both extra items and increments for cues serve to reduce sampling probabilities for critical words. On the one hand, SAMS provides a detailed and accurate account of the list length effect (see the fit to Roberts' data); on the other hand, SAMS does not need to posit a lengthening of the list or incrementing (see Figure 9) to explain the effects of part-list cuing. Nevertheless, a list length factor could quite possibly be contributing something to the observed effect.

Why have previous associative network models (e.g., Anderson, 1972; Anderson & Bower, 1973) had difficulty predicting the part-list cuing effect? Although they do include item search along associative pathways, they fail to utilize a random search. Thus, many items in the control condition may not be reached because entry points to memory are limited. Furthermore, all items that can be reached (within certain limits) from the list cues will be retrieved, and there is no cost involved in doing so. There is no build-up of failures nor a stopping rule based on such a build-up. These assumptions are

of course sensible when a directed, nonrandom search is postulated, since the search is assumed to proceed serially with no real provision for resampling of previously found items. Thus search simply proceeds until all recallable items have been found.

The present theory, on the other hand, does incorporate the notion of costs involved in retrieval. Thus the subjects are given a reasonable basis for cessation of search, even though potentially recallable items remain in memory. We have shown that a sensible stopping rule applied equally to both conditions, no matter how strict the stopping criterion, can give rise to a cue condition inferiority. The basis for SAMS' solution for the part-list cuing puzzle is the combination of an associative network (as represented in a retrieval structure) with a cue-dependent probabilistic sampling approach.

IV. The SAM Theory

Guiding Principles

A number of general considerations are precursors of the SAM theory.

1. Long-term memory is effectively permanent, with additions allowed, but not deletions. As a corollary, forgetting is a result of retrieval failure.

2. Long-term memory is a richly interconnected network, with numerous levels, stratifications, categories, and trees, that contains varieties of relationships, schemas, frames, and associations. Roughly speaking, all elements of memory are connected to all others, directly or indirectly (though perhaps quite weakly).

3. Memory retrieval is cue dependent. What image is elicited from memory is determined by the probe cue utilized at that moment (Tulving, 1974).

4. Memory retrieval is noisy and hence probabilistic. A given set of cues has some chance of eliciting from long-term memory any associated image (though for many images the probabilities may be extremely small). Furthermore, successive use of the same probe cues may well result in the elicitation from memory of different images.

5. Temporal-contextual information is of fundamental importance; its role in storage

and retrieval must be delimited carefully and explicitly.

Overview

SAM assumes a partition of memory into unitized images. These images are the objects in memory that may be sampled during a memory search. Images may be quite complex information structures and may overlap considerably with one another. An image is defined by a property of unitization: At a given stage of the memory search, the information recoverable from memory is limited to that in some one image.

The basis for retrieval is assumed to be the strength of associative relationships between probe cues and memory images. These are described within a retrieval structure that is a matrix of retrieval strengths from each possible cue to each possible image. The strengths determine the tendency for a given probe cue to elicit a given image when long-term memory is probed with a given set of cues. A retrieval structure is not a storage structure; it is usually much simpler. However, the retrieval structure is designed to capture those aspects of the storage structure that are important for retrieval.

Many images, especially those of importance in laboratory tasks, are organized with respect to temporal-contextual information. In such tasks, contextual information will always be one of the cues used to probe memory. Other cues, such as item and category names, may be used as cues in addition to context. A quantitative ratio rule is posited in SAM that determines the probabilities of sampling each image, given any number of cues in the probe.

When an image is sampled, some of the information in the image becomes available to the subject for evaluation and decision making. The proportion of such information recovered is determined by the retrieval strengths between cues and image.

Retrieval is assumed to consist of a series of search and recovery operations organized by a retrieval plan. The plan may be altered as search proceeds. It is used to determine what cues are utilized at each stage of the search and to make certain decisions, such as when to terminate the search. Retrieval

also includes a rapid initial phase, called sensory coding, that takes place when new information is presented. This phase is largely automatic, usually concludes within several hundred milliseconds, and generally results in the accurate retrieval of features that are representative of the input in any context (including low level codes and a certain amount of semantic or type information). Our main concern in this article is with memory search following the sensory coding phase.

The Units of Long-Term Memory

The conception of long-term memory as a richly interconnected system raises immediate problems for memory theorists. It would be impractical and inadvisable to think of memory as a single complex object. On the other hand, the boundaries of "objects" in an interconnected system are bound to be imprecise. It is common, for example, to think of a word as an object when analyzing traditional memory paradigms involving word lists. However, such an object consists of a constellation of related concepts and elements including everything from shapes, sounds, phonemes, and letters, through parts of speech, synonyms, and meaning, through codes, sentences, and conversations and stories that might have contained that word, to the temporal-contextual setting in which the word was presented and possibly other words presented at about the same time. All of these may form part of the "object" designated as a word. Yet all of the elements that make up an object will themselves be associated, to some degree, to other elements not in the object. In such conditions, an object's boundaries are bound to be somewhat fuzzy, at best.

A second problem that makes it difficult to partition memory into objects is "level." If subjects are presented with and asked to remember letters from the alphabet, then letters might be the basic level of the information in each object. In other tasks, single objects could comprise words, sentences, paragraphs, stories, or pictures.

A third problem depends on the structure of memory. Suppose for example, that a portion of memory is organized as a hierarchical

tree. Should the entire tree be treated as a single object, should the individual nodes be defined as the objects, or should some intermediate partitioning of the tree be defined as the object structure? The answer is far from obvious and may even differ from one task to another.

Despite these problems, it is both necessary and desirable to partition memory into objects; in SAM this is done through the concept of *unitization*. It is possible to distinguish objects with imprecise boundaries by assuming that interconnections between elements will be stronger and more numerous within one object than between objects. When retrieval from memory occurs, a set of elements may be activated, which is thought of as entry into short-term store. Short-term store is limited in capacity (Shiffrin, 1976), and unitized sets of information presumably require less capacity. Thus it may be easy to maintain elements of one object in an active state, but difficult to maintain the same number of elements if they are spread among several objects. We assume for this reason that retrieval operates on one object at a time. The information recovered from long-term store during one loop of the search process will be a subset of the information in one object. (How the object is selected will be discussed later.) The unitized objects of long-term memory are henceforth termed *images* (no visual representation is implied).

How does unitization occur? The most important factor is the encoding and rehearsal process. The information and structure that are simultaneously in short-term store undergoing coding and rehearsal, whether composed mainly of letter, word, sentence, or story information, tend to be stored as a unitized entity. Of course, simultaneous presence in a rehearsal buffer is not sufficient for a unitized entity to be formed. The nature and amount of coding and rehearsal will help to determine whether a unitized image develops. A unitized memory object will form most easily when the inputs to short-term storage consists of already unitized memory entities. This will be the case with words, say, or nursery rhymes. In such cases, only the present context would have to be combined with the previous im-

ages in order to produce a unitized image for the current situation. These assumptions imply that there is a good chance that new images will form when the same item is repeated in different contexts or is repeated in a similar context but coded in a new way.

It should be kept in mind that the various images in memory will, in general, overlap considerably, in the sense that the subcomponents of one image may be the subcomponents of other images as well. For example, sentences might be stored as the images in memory, even though the same words may appear in more than one sentence. In such cases, the images as a whole are quite distinct, even if the subcomponents are not. Note that one of the most important subcomponents that helps to distinguish images from each other is the temporal-contextual information that is stored within each image.

When an image is formed by the addition of temporal-contextual information to pre-existing images, it can be described as episodic in Tulving's (1972) terms or as a token in Anderson and Bower's (1973) terms. Such images will tend to be retrieved selectively when one of the cues is an appropriate context cue. Note that the formation of the new image (e.g., *horse + context*) does not remove or replace the previous image (*horse*) from memory. Whether that previous image was episodic (*horse + some previous context*) or semantic (*horse + no particular context*), it will still remain in memory. Such reasoning is consistent with our general view that new inputs to memory are additive but not subtractive.

Of course, images can be formed that have no particular contextual basis—these are usually called semantic images or types. How are they formed? One possibility arises when the contextual coding is weak or ineffective. For example, in tasks where meaning rather than setting is emphasized, the supplementary information stored with the word in the new image might be semantic in nature rather than time or situation specific. In addition, features that almost always are present when an input is repeated in different contexts will tend to become part of the feature set that is automatically retrieved during sensory coding. Sensory coding will be discussed later.

Are the objects of memory, once formed, fixed and inviolable forever? One alternative answer is that objects are partitioned with respect to a given set of probe cues used in retrieval; different probe cues could result in a change in the level and boundaries of the objects. For example, stories might be presented and stored as images, but a later recognition test for individual words might lead the subject to desire to tap memory at the word level, rather than the story level. It is possible that the partitioning of memory into objects depends to some degree on the choice of probe cues. However, the tasks and situations to which we are currently applying SAM (or SAMS) do not seem to require this assumption. In any event, selection of images at different levels could be accomplished by judicious selection of cues.

It should be understood that there may be only a slight and subtle distinction between, say, two words plus context as a unitized entity and two words, each with context, as separate entities connected by an association. There are differences, however, that are potentially testable in the context of our sampling and recovery assumptions. After sampling of a unitized two-word image, all of the available information from both words would be used for decision making, with no need for further search. On the other hand, when one of two separate word images is sampled, only the information in that image is recovered. Recovery of information from the other image would require further searching (perhaps with the first word as an additional cue), which might not succeed even when the between-image association is quite strong. SAMS assumed for simplicity that all images were single-word units (with appropriate interconnections) and predictions were quite accurate. In other tasks, the assumption of multiword units might prove necessary.

In summary, the interconnected long-term system is assumed to be constituted of a class of relatively unitized images. These images may be interconnected in extremely complex ways: in networks, hierarchies, and with various types of relations. Furthermore, a given image will tend to have a complex infrastructure. An image has the following important property: When sampled in a mem-

ory search, recovery of the available proportion of its informational content occurs directly, without a need for further memory search with the same probe cues.

Storage Structures Versus Retrieval Structures

Because our main interest lies in the development of a retrieval theory, very few assumptions will be stated concerning the interimage structure. Our aim is to develop an approach to retrieval that can be used for as wide a class of storage structures as possible. To do this, we introduce the concept of a *retrieval structure* consisting of retrieval strengths between the possible cues (images and their subcomponents) and the images in memory. The strength determines the tendency for a given cue to sample (i.e., elicit) a given image and should be larger the longer that the cue and image were coded or rehearsed during storage. This strength does not necessarily have to depend on the type of interimage connection that was coded. For example, *daisy* and *rose* presented in a list might be coded through the relation *both are flowers*, through visual images of each in a vase, through connection to a superordinate node or control element such as *flowers in list*, or indirectly through a chain of associations between intermediate nodes or intermediate list words. It is not necessary that *rose* as a cue be more likely to elicit the image of *daisy* in one of these cases than the other. In each case, the retrieval strength (and hence the sampling tendency) would be determined by factors such as the effort and duration expended in coding and the type of coding carried out (e.g., rote rehearsal vs. mnemonics). For the purposes of sampling, then, the memory structure need not be taken into account explicitly (see Anderson, 1976, p. 154), as long as the images and cues to be placed in the retrieval structure are chosen judiciously (e.g., on the basis of the subject's coding strategies), and the cue-to-image strengths are chosen sensibly (e.g., on the basis of the effort, type, and duration of coding).

The foregoing discussion should not be taken to imply that the interimage structure is irrelevant for retrieval. The interimage

retrieval strengths should reflect the actual interimage structure closely, so that even sampling could be said to conform to the stored structure in this sense. More important, the details of the between-image relations, whether relations, labeled associations, or anything else (see Anderson & Bower, 1973; Kintsch, 1974; Norman & Rumelhart, 1975), are part of the information contained in each sampled image. This information, when recovered, will help to determine various decisions, evaluations, and reconstructions; the recovered features could also be used as additional cues during later stages of the memory search. A good example was the application of SAMS to categorized free recall: Category membership was stored in each image, and category cues were used in the cue sets.

Long-Term Retrieval

A SAM retrieval theory is organized by a *retrieval plan* that governs the successive sampling and recovery operations from a retrieval structure. Learning during retrieval is termed *incrementing*, and the initial, rapid retrieval when new information is input is termed *sensory coding*. These various components are discussed in the following sections. Except where indicated to the contrary, it will be assumed for simplicity that the partitioning of long-term store into images does not change with different cues.

The retrieval plan. All decisions and control processes involved in retrieval are subsumed in the retrieval plan. The organization of retrieval is an extension of that proposed by Shiffrin (1970). Retrieval is a memory search proceeding in series of discrete steps. Each step involves a probe of long-term store by one or more cues, which results in a briefly activated set of information, which is followed by a selection, or sample, of an image from that set. The substages within any one step are depicted in Figure 2. Retrieval begins with some question the subject needs to answer regarding the contents of long-term store. This may be as simple as: What is another word on the list most recently presented? In the most general case, a retrieval plan is generated to guide the search for the answer. Initially, the plan may

be somewhat vague by intention, in the hope that later phases of the search will be guided by information located in earlier phases. The plan includes such things as an initial decision whether to search long-term store; how to search (for instance, in a temporal order or by an alphabetic strategy); how to choose probe cues (for instance, should recalled information be used as probe cues?); what combinations of probe cues should be employed and with what weights; whether to employ the same probe cues on successive loops of the search or whether to alter the cues; whether to search first for preliminary cues to guide later search; and how long to search (i.e., how many loops of the search process should be carried out). The retrieval plan itself is constructed on the basis of the information in the test query, the information currently available in short-term memory, and information retrieved from long-term memory; this information retrieved from long-term store may be concerned with search plans, previous successful plans in similar situations, and so forth. Once constructed, the retrieval plan is also stored in LTS. We assume in the context of the tasks discussed in this article that retrieval of the plan itself is easy and accurate, but there could be cases where this assumption is violated. Simple retrieval plans that might plausibly be used in free-recall tasks were discussed in Sections I and III. In general, however, little is known about retrieval plans, and their exploration may well become an active research area in the near future. (See Williams, 1978, for a discussion of retrieval plans.)

Next, on the basis of the retrieval plan, the subject assembles probe cues to be used in retrieval. These cues may include: (a) information the subject has about the context at the time of study, (b) context representative of the moment of test (although these cues may not be useful or desired), (c) information from the test question, (d) information retrieved earlier in the search, and (e) information generated during construction of the retrieval plan.

After probing LTS, the cumulative activation might be utilized as a basis for a recognition decision. If search does not stop at this point, then the information recovered

from the sampled image is evaluated. The decisions involved depend on the task but might include the following: Does the image derive from an item on the presentation list? What is the name encoded in the image? Is the encoded name an appropriate response to the cue? Does the encoded name match the test cue? Does the encoded information hold the proper relationship to the probe information? Is any of the recovered information useful for later phases of the search?

At this point the subject may or may not decide to emit a response. Then a decision must be made whether to continue the memory search. Presumably the decision to terminate is based on the previous successes or failures up to that point. For example, in free recall a decision to stop might occur when m successive searches occur with no new words recalled. If the search continues, the search loops back to the retrieval plan where the next probe set is chosen, and so forth.

The changes in probe cues as the search continues are quite important. Although sampling is a random process, it is easy to misinterpret such a statement and ascribe more randomness to the retrieval system than is, in fact, present. The strengths may be such that one image is far more likely to be selected than any other. Even more important, the subject can control the search by changing the probe cues as needed. Systematic changes in probe cues may not lead to rapid conclusion of a memory search but may be quite effective nonetheless. For example, when a subject is asked to recall a United States city starting with the letter X ,² one strategy would involve searching with state-name cues, one at a time, generated systematically in geographic fashion.

The retrieval structure. A retrieval structure is a matrix of the strengths between the possible probe cues and the various memory images. It is a generalization of the strength matrix shown in Figure 1. The cues may be labeled Q_1, \dots, Q_m , the images I_1, \dots, I_m , the strengths $S(Q_i, I_j)$. In theory, all possible cues and images are contained in the structure, but in practice, most of the structure is irrelevant. Typically, therefore, the matrix is restricted to images and cues

² Xenia, Ohio.

that are task specific. In SAMS, for example, the images and cues were restricted to context, list words, and category names.

The cues and the images in a retrieval structure are generally quite distinct even when they each encode nominally equivalent information (e.g., the same word). For example, suppose *horse* has been placed in long-term memory at the time of study, and then a recognition test is given with *horse* as the test item. We argue that the image and the test cue are quite distinct, the image consisting of *horse at study + study context* and the test cues consisting of *horse at test + test context* (the encodings of *horse* may differ in the two instances, perhaps *rocking* vs. *running*). It is assumed that the long-term memory image and the probe cues are separate and identifiable entities, so that if the memory image is sampled, it may be evaluated alone or compared with the probe cues. However, these two entities will usually be strongly associated due to their large pool of common information; it is this fact that makes it likely that the cue *horse* will cause the image *horse* to be sampled. When an image and cue encode the same nominal information, the cue is termed a *test analogue* of the image.

Note that the similarity of an image to its test analogue used as a cue can be greater than is the case in the example above, if the cue is itself retrieved from memory during a previous loop of the search. The reason is clear: The context in the cue in this case would be more similar to that at the time of storage, since some of it would have been retrieved. Such a cue might also have higher strengths to other images in memory.

Such reasoning suggests that the retrieval structure should contain two types of cues whenever cues are provided by the experimenter—one type consisting of retrieved cues, the other, provided cues. Although such reasoning is correct in principle, it may be excessively pedantic in practice. Presumably, a rational subject accumulates and utilizes all relevant contextual information as he proceeds in the search. As a result, with the possible exception of the first few samples from memory, the difference between provided cues and retrieval cues may be small.

To generate a retrieval structure for a particular setting (rather than the theoretically infinite-size structure), one would begin by choosing an appropriate coding model to apply during storage. The storage assumptions determine what type of images form and with what probabilities. The cues are generally chosen to represent the test analogues of the set of images that have been formed, along with the natural subcomponents of those images (particularly the context component). Additional cues are placed in the matrix if they are provided by the experimenter. Thus, for example, if sentences are presented, the images will all contain context information; some might also contain sentence concepts and others might instead contain fragments of these, including perhaps single-word concepts. The cues will then consist of the test analogues of those images and their subcomponents.

The strengths in the retrieval structure depend on the assumed storage process, on preexisting associations and relations in force before storage and at test, and on changes in context between storage and test. The strength determines the tendency with which a given cue tends to elicit a given image during sampling (and also determines the amount of information recovered from the infrastructure of a given image). The more the information in a cue matches the stored information in an image, the greater is the strength. Thus the longer two words are rehearsed together, the greater the strength between one of the words when used as a cue and the image of the other word. On the other hand, as the time between study and test increases, or as the similarity of the context at study to that at test decreases, the between-word strength should decrease. For list words that were not rehearsed together, the strength (though low) will depend to a greater extent on extra-experimental factors and should be affected less by delay and context change.

In summary, the retrieval structure is not meant to be a copy of the stored memory structure. Rather, the retrieval structure may be thought of as a scaled simplification that captures aspects of the storage structure that are relevant for cue dependent sampling.

The sampling process. All SAM retrieval theories assume the same sampling process. In words, an image has a probability of being sampled that is determined by the associative strength relating the set of probe cues to the image in comparison with the strengths relating all other images to the set of probe cues. More precisely, given a retrieval structure, the probability of sampling Image i , given cues Q_1, Q_2, \dots, Q_m in the probe set, is given by

$$P_S(I_i|Q_1, Q_2, \dots, Q_m) = \frac{\prod_{j=1}^m S(Q_j, I_i)^{w_j}}{\sum_{\substack{k=\text{all} \\ \text{images}}} \prod_{j=1}^m S(Q_j, I_k)^{w_j}}. \quad (4)$$

A few words of explanation are useful to aid understanding of Equation 4. First, note that this equation is an extension of the Luce (1959) ratio rule (most often applied to choice behavior). This is most evident when there is only a single cue and $w = 1.0$, in which case Equation 4 becomes $S_T(Q, I_i) / \sum_k S_T(Q, I_k)$. This special case was also the basis for the Shiffrin (1970) search theory.

In Equation 4, the w_j are weights representing the saliency of the j th cue, or the attention given the j th cue.³ In SAMS, the w_j were all set to 1.0, regardless of the number of cues (though no more than 3 cues were ever used together). Because STS is limited in capacity, there surely is a limit on the number of probe cues that can be utilized at one time, and a limit on the amount of attention that can be distributed among the various cues. Such limitations could be captured in any of several restrictions placed on the w_j . One possibility is a limit, W , on the sum of the weights:

$$\sum_{j=1}^m w_j \leq W, \quad \text{for all } m. \quad (5)$$

The SAMS model would be consistent with this restriction if W were at least 3. A somewhat different assumption posits a fixed capacity:

$$\sum_{j=1}^m w_j = W, \quad \text{for all } m. \quad (6)$$

The possibilities of Equation 6 and of the use of nonunitary weights are of more than academic interest. If weights are always equal to 1.0, then the similarity or overlap among cues is not properly taken into account. To see this, consider an example of extreme similarity among two cues: The cues are identical. Substitution into Equation 4 shows that the use of two identical cues, say A and A , does not give rise to the same sampling probabilities as the use of A alone. This logical difficulty could be handled by relaxing the assumption that all weights are unitary. For example, the weight given to a cue could be reduced by a factor representing the cue's overlap with other cues. Alternatively, the problem could be handled by adopting Equation 6. In this case, A and A together would give rise to the same sampling probabilities as A alone. The SAMS simulation worked successfully despite letting all weights be unitary. We suspect the reason lies in the choice of cue sets: They were always made up of dissimilar cues.⁴

Aside from the weights, the sampling rule assumes that the products of the strengths of the cues to an image are used in the ratio rule. Such a product rule allows the search to be focused on images that are strongly connected to all the cues, rather than just one of the cues. One might think that it would generally be advantageous to combine as many cues as possible in each probe set to focus the search as narrowly as possible. Even when the weights are all 1.0, however, so that recovery improves with additional cues, this is not necessarily the case. One problem is the tendency for sampling one of the images of the cues themselves; this tendency could well be disadvantageous in certain recall tasks. For example, in the SAMS applications to part-list cuing, we saw that one factor inhibiting free recall (for both conditions) was excessive reliance on word-plus-context cues, compared with the context cue alone. Furthermore, there are var-

³ Note that our use of a multiplicative rule requires the weights to appear as exponents. If they were multiplicative factors, they would simply cancel out of the sampling equation.

⁴ The assumption of Equation 6 would change the SAMS simulation. We have not yet explored this alternative model and do not know whether the predictions would change in important ways.

ious dangers involved in focusing a search too narrowly, since a desired image may be missed. Finally, the use of multiple simultaneous cues may be restricted by capacity limitations such as those described by Equations 5 and 6.

Most of the images in long-term memory have such low retrieval strengths to the cues that their sampling probabilities will be vanishingly small. The relatively small set of images with nonnegligible sampling probabilities is called the *search set*. It is therefore convenient (especially when incorporating the model in a computer simulation) to separate the sampling phase into two parts: first, a restriction to the search set; second, an appropriate probabilistic choice from the search set. The choice of search set is generally determined by task considerations. For example, if a subject is asked to recall a just-presented list, the search set might be assumed to consist of the images of all the words in that list (or perhaps of all the words in the session, if intrusion predictions are desired). Of course the model should always specify a search set containing the images that might be sampled. If there is reason to believe that images are sampled that are not in the search set specified by the model, and relevant predictions are desired, the search set should be expanded to include the additional images.

Recovery. When long-term memory is probed by a set of cues, a fairly large collection of information in many images is very briefly activated. The total amount of such activation (and the relative activation of different images) is determined by the strengths in the retrieval structure for the cues that are used. The activation of this information is equivalent to the placement of the information into short-term store (STS). Because STS is limited in capacity, it is incapable of retaining very much of the activated information for more than a short time, perhaps a few hundred milliseconds. SAM assumes, therefore, that information from just a single image is retained long enough that the subject can make a detailed evaluation during that loop of the search process.

For a given set of cues, the recovery assumptions first must determine the amount and type of initial activation. Then, when an image is sampled, the recovery assumptions

must determine the amount and type of information that can be extracted from that image. The total amount of initial activation is probably proportional to the denominator in Equation 4; it can be viewed as a feeling of familiarity and could conceivably be used to make a recognition judgment without further search. Detailed information is not available, however, until recovery from a sampled image occurs. Thus, in tasks such as recall, in which the details of the information are a prerequisite for a response, only the recovery from sampled images is relevant.

The process of recovery is assumed to be noisy and imperfect, so that not all of the elements of an image will be activated. In general, the stronger the retrieval strength between the sampled image and the probe cues, the larger will be the proportion of image elements that will be recovered and made available to the subject's evaluation and decision making mechanisms. In SAMS, the same retrieval strengths that determine sampling probabilities are used to determine the proportion of recovered elements. In general, however, these strengths might not be equal; their relationship will depend on the response requirements of the task.

Once information is recovered, it is evaluated and used in decision making and to generate responses. In practice, it is often convenient to combine the recovery process and the subsequent decisions into a single equation. Consider, for example, a task in which the sampled images are words. In such a case we propose the following equation to give the probability that a sampled image, i (that has not previously been sampled), can give rise to an accurate production of the name of the encoded word, when Q_1, \dots, Q_m are the cues:

$$P_R(I_i | Q_1, Q_2, \dots, Q_m) = 1 - \exp\left\{-\sum_{j=1}^m w_j S(Q_j, I_i)\right\}. \quad (7)$$

This equation was utilized in SAMS, with the w_j all set to 1.0.

The form of this equation is somewhat arbitrary mathematically, though it does capture a number of features we consider desirable for a recovery rule in this case.

First, the stronger the strength is to any one cue and the stronger the summed strengths are to all cues, the more likely recall is. Second, the larger a cue weight is, the more the strength to that cue will affect recall. Note that this production rule obeys an additive rather than multiplicative rule, so that recall will be high if even one weighted strength is high. Third, the probabilities will range from 0 to 1 as the sum of the strengths ranges from 0 to ∞ .

It should be noted well that in this system, the recovery probabilities are determined by strengths from cues to image, not by image strength. Image strength per se does not in fact exist in the system. Undoubtedly, the names of common words are well encoded in numerous images in long-term memory. Whether the name can be generated in response to the cues (including context) is the only relevant question.

Consider next how recovery attempts should be related when the same image is sampled several times during a search. Shiffrin (1970) implicitly assumed independence of the information recovered in successive attempts, but such an assumption is difficult to defend. We propose here an alternative view—that over the short time span normally taken by recall, the information recoverable from a given image does not change appreciably when the same cues are used to resample that image. However, if even one new cue is involved in the resampling, then substantial amounts of new information might be recovered.

The implications of these assumptions for recall tasks are straightforward. They were instantiated in the SAMS simulation in the following manner. It was assumed that successful recall of the word encoded in an image is always followed by successful recall from that image, regardless of cue set. Failure to recall was assumed to be followed by other failures to recall, unless at least one new cue for that image is in the cue set; in this case, a new independent chance at recall occurs.

It will not escape the reader's attention that the recall and recovery rule of Equation 7, which was used in SAMS, makes no use of the fine structure of the association by which a cue and image might be related.

Only the strength value is utilized.⁵ Such a view can be defended, perhaps, in consideration of the structureless tasks of free recall to which SAMS is applied. For tasks such as sentence and paragraph memory, it would not be surprising if the recovery and evaluation rules required explicit consideration of the type of relational information between cues and images (which is assumed in SAM to be stored in each image). The point is the following: Although SAM does not take the nature of an association into account during sampling (beyond the degree captured by a strength value), there is no prohibition of recovering and using such information after sampling.

Incrementing. Many results suggest the necessity of a storage process that operates during the course of retrieval (see Sections I and III and Raaijmakers & Shiffrin, 1980). One type of storage is particularly easy to justify, since we have assumed that storage occurs when information is rehearsed together in STS. Since information recovered from a sampled image enters STS, where it is considered in relation to the probe cues already there, it seems likely in this case that the retrieval strength between the probe cues and the sampled image should be increased. In a recall situation, in which the name is required, it seems likely that the major increase in strength would occur after successful recall takes place, with little or no increase after a failure to recall.

The process of increasing cue-to-image strengths during retrieval is called *incrementing*. Assuming that significant incrementing occurs only after successful recovery, it remains to be established how much incrementing occurs when the same image is sampled successfully a number of times in a row with the same cues. It hardly seems likely that incrementing would occur unattenuated after each successive sample of the same image. (Since each increment increases the sampling probability of that image in the

⁵ Actually, the use of category cues in the case of categorized free recall and the assumption of the recovery of category information from a recalled image were mild concessions in SAMS to the necessity of taking the type of relationship into account. This does not alter the basic point made in the text.

future, the selection probability for some one image could eventually become almost unity.) We propose, therefore, that the major increase in strength occurs after the first successful recovery, with increases becoming smaller thereafter. In SAMS, in fact, the simplifying assumption was made that incrementing between a given set of probe cues and a given image occurs only after the first successful recovery. These assumptions would of course be altered if the subject were led or instructed to give special attention and rehearsal to recovered items.

Long-Term Forgetting as Retrieval Failure

There are two basic reasons why an image may be retrieved better at Time A than at Time B. First, the cues utilized at Time A may be more strongly associated to that image than those used at Time B. Second, the strength and number of other images associated to the cues (even if the cues are the same) may be greater at Time B than at Time A. Everything else being equal, an increase of cue strength to a given image increases both sampling and recovery probabilities (see Equations 4 and 7). On the other hand, for fixed cue-to-image strength, an increase in the strengths of cues to other images will reduce the sampling probabilities (though leaving recovery unaffected).

The increase in the strengths of cues to other images tends to be an inevitable consequence of new learning. This new learning will not necessarily lead to forgetting, however. The new information might be organized together or integrated with the old image so strongly that a larger image is formed that contains both the new and old information. Alternatively, the new information might form a new image that is tightly associated to the old image. Retrieval of one of these images could result in that image's being used as a cue and eliciting the other image. In this case forgetting is prevented by an appropriate switch of probe cues. Conversely, when the cues are not changed, forgetting may occur. This leads to the general principle that forgetting due to new learning occurs when the same cue is utilized in an attempt to locate one image among an in-

creasing number of other images. On the other hand, the subject may change the cue selection during the search so that each cue is related to a subset of the increasing number of images; in this event forgetting may be lessened or even reversed.

The decrease in the strengths of association of cues to image can be the result of several factors, chief of which is the change of context over time (e.g., see Bower, 1972; Estes, 1955). The context at the time of storage of an image makes the best retrieval cue for that image. However, the context cue used at testing may consist largely of the context information available only at that time. Since the test context usually differs from the storage context by a greater amount as time between storage and test increases, the retrieval strength is reduced.

The role of contextual factors in retrieval and forgetting is very important, and it is an essential component of SAM. The temporal-contextual features include incidental information from the sensory environment and the subject's long-term store that happens to be present in STS at the time of a storage event. They might include the location, the temperature, the time of day, recent events, and the subject's physical state, feelings, emotions, and recent thoughts (Bower, Monteiro, & Gilligan, 1978; Smith, Glenberg, & Bjork, 1978). In SAM, such context features are always one of the probe cues, either by intent of the subject or by accident. Presumably, the subject can, through attention, vary the weight assigned to this context cue, but such information is always present in STS and always plays at least a small role as a retrieval cue. Whenever possible, of course, a knowledgeable subject tries to reinstate in STS as far as possible the contextual cues that were present at the time the to-be-recalled image was stored (Smith, 1979).

The first forgetting factor, in which one set of cues is connected to an increasing number of images, was utilized extensively by Shiffrin (1970) to explain list length effects in free recall. Since that time, list length effects (sometimes called *fan* effects) have been observed in many settings. In our own research, we have observed such effects in paired-associate and recognition paradigms, as well as in free-recall studies. Muel-

ler and Watkins (1977) have referred to this forgetting factor as *cue overload*. The SAM theory may be said to provide a theoretical basis for the cue overload principle. It is interesting to note that Mueller and Watkins used the cue overload principle to explain the part-list cuing effect. As described earlier in this article, the SAMS model explains the part-list cuing effect without any recourse to a cue overload factor. Nevertheless, the cue overload principle is, in SAM, one of the fundamental bases of retrieval failure.

The second basic forgetting principle, context change, or more generally, reductions in cue-to-image retrieval strength, is used to explain the deleterious consequences of delay of test and helps explain the difficulties caused by the presentation of interfering materials during the study-test interval. An exposition of this principle would require a model of context change and is well beyond the scope of this article.

Sensory Coding

When a stimulus is first presented to the subject, it begins undergoing sensory analysis. Even the very lowest, initial stages of sensory analysis are properly viewed as involving retrieval from long-term memory. At each stage of analysis, the previously retrieved features plus context already in STS determine the next features to be retrieved. When the stimuli are highly discriminable, these stages of sensory coding are rapid, accurate, and automatic. At the higher levels of analysis, the meaning of the input is retrieved along with, perhaps, particular contextual images containing the input. Although it is difficult to draw a dividing line at any one point in this stream of processing, we lump together the relatively noncontextual retrievals and call the results *sensory coding*. These retrieved features are then viewed as the first item cue for the subsequent memory search.

Sensory coding typically concludes in several hundred milliseconds and is quite automatic, differing in both respects from the subsequent memory search. Nevertheless, it is possible that the same basic retrieval

mechanisms (e.g., Equations 4 and 7) operate during this initial phase of retrieval. Such a possibility could be explored within the context of recognition and detection paradigms. Of course, if the inputs are degraded so that accuracy of coding drops and confusions result, the situation is even closer to that in the memory search domain. Perhaps threshold detection could prove consistent with a search model positing a single sample from memory. Such a possibility is not entirely implausible, since one of the successful models for sensory confusions is a modified Luce choice model that is formally identical to Equation 4 (Luce, 1959; Luce, Bush, & Galanter, 1963).

Applications of SAM to Other Paradigms.

Although it is not possible in this article to describe in any detail the application of SAM models to other paradigms (see Raaijmakers & Shiffrin, 1980), a brief discussion is quite useful, especially since very few changes in the SAMS simulation are needed to fit the results from certain standard paradigms.

Free recall. No changes in the structure of SAMS should be necessary to fit the results of standard, verbal, free-recall paradigms. Cumulative responses as a function of time and interresponse times are already inherent in SAMS, since each loop of the search may be assumed to occupy one unit of time. Output rates slow down, since resampling of previously sampled images generally increases at the expense of new sampling as search continues. The resampling tendency is increased even further by incrementing (one tends to sample items previously recalled). It would be possible for a subject to reduce the resampling tendency considerably by using a systematic strategy to generate and change cues. For example, a search based on first letters could proceed systematically through the alphabet. Apparently subjects seldom use such strategies.

Sometimes subjects are asked to recall the same list several times in succession with no new presentations. About the only addition to SAMS that is required to deal with this situation is a decision concerning whether

additional incrementing should be allowed when the same word is recalled during successive retrieval periods.

SAMS makes no provisions for confusions or intrusions in free recall. Intrusions of words in the session but not on the current list can be handled by expanding the retrieval structure to include such images (with low strengths to current list cues, of course). Intrusions of nonpresented words can be handled in part by allowing for imperfect recovery of presented words, that is, confusions. Another possibility would involve expanding the strength matrix to include images of nonpresented words (with very low strengths to current list cues). Note that an implicit part of the recovery process in SAMS is a judgment that the sampled word was in fact on the current list. If false alarms could occur, then images of words not on the current list could be mistaken for current list words, and intrusions would occur (as opposed to confusions).

The application of SAMS to categorized free recall has already been described in this article, albeit quite briefly. No important change in the SAMS model was required, except for the inclusion of a category name cue in the retrieval structure and the assumed use of such a cue during retrieval.

Paired-associate paradigms. Experiments can be carried out in which items are presented only once prior to test but in which the items are grouped at input. Then testing could be done by a free-recall technique or by presenting one or more members of each group as cues and requesting recall of the remaining members of the group. If the grouping is by pairs, and one member of a pair is provided as a cue at test, the paradigm may be termed *single-list paired-associate* recall. It should be clear that SAMS is structured to deal with such a paradigm, since cued testing is already a fundamental part of the model and is used extensively in the applications described earlier in this article.

We have found that the results of such studies can be predicted quite well by SAMS, even though the images are all assumed to consist of single words plus context. One might ask whether there are data requiring multiword images to be added to the re-

trieval structure. We have not found this necessary in paired-associate settings but have found such assumptions helpful in paradigms involving three-word groupings and cued testing. We suspect that such assumptions might also be useful in studies involving sentence presentation. In fact, Jones (1976) has fitted a variety of such data with a fragment model that incorporates a central assumption that stored units can consist of sentence fragments ranging in size up to full sentences. Although the SAM theory differs from the fragment model in many respects, an assumption of multiword images in SAM might incorporate many of the same desirable properties that are captured by fragments in Jones' model.

Learning paradigms. There is a vast literature on verbal learning experiments that involve multiple presentations of items or lists. A variety of changes would take place during learning in such studies according to a SAM theory. These include the development of the following: multi-item images, a complex interitem structure, an elaborate dependence on context, and a systematic retrieval strategy. These complexities take discussion of such paradigms beyond the scope of this article.

Recognition paradigms. Recognition poses particular problems for search theories. Simple theories require an extended search without success before an accurate response can be given that a test item is not on the presentation list, but if this were correct, the possibility of rapid, accurate, *no* responses would be eliminated. The problem is not restricted to stochastic search theories, but this is not the place to discuss the accuracy of such reasoning or the relevant data (see Ratcliff & Murdock, 1976). Instead we propose one way in which rapid recognition responses could be produced within a SAM framework.

It is proposed that the subject has access to the sum of strengths that appears in the denominator of the sampling equation (Equation 4). The details of the activated information are probably not available at this stage. A value proportional to the sum of strengths could be described as a "feeling of familiarity." Within a given task, criteria

could be set such that a familiarity value above the upper criterion could be used as a basis for a positive recognition response, and a familiarity value below the lower criterion could be used as a basis for a negative recognition response. A familiarity value between the criteria would be ignored, and a normal search would proceed thereafter. Such a model is similar to that of Atkinson and Juola (1973).

Ignoring the possibility of a recognition response based on familiarity, recognition in a SAM theory proceeds just as in recall settings. The sampling is identical, but the image sought is one that contains the probe cue itself. In addition, after information is recovered from a sampled image, several judgments must be made. A *match* judgment determines whether the information in the sampled image could be an encoding of the test item. A *context* judgment determines whether the sampled image was on the relevant presentation list. Note that these judgments must depend on the quality of the recovered information; strength alone is not sufficient. For example, a change in context reduces the cue-set-to-image strength, and hence less of the information in the sampled image is recovered. But will a judgment of a match be increased or decreased as a result? More detailed knowledge of recovery is needed before this question can be answered.

Final Comments

We have presented in some detail the underlying basis for a new retrieval theory that assumes probabilistic, cue-dependent sampling from a retrieval structure representing an associative network. The storage processes, memory representations, and retrieval mechanisms have all been laid out, although the greatest emphasis has been directed toward the search and retrieval operations.

A specific model, SAMS, was developed for free recall. Parameters were estimated so that the model would fit free-recall data involving variations in list length and presentation time. SAMS was used, essentially intact, to fit a variety of results from the part-list cuing paradigm. Although the model contains many processes and parameters, the predictions of the model are largely inde-

pendent of the process assumptions and choices of parameter values. In fact, the predictions graphed were based on parameter values independently generated in fitting Roberts' (1972) free-recall data. The predictions are impressive for this reason and because they are based on a model making extensive use of interitem associations in retrieval—a possibility thought by previous investigators to be ruled out by the data.

There are a number of strengths of SAMS and also of the general theory, SAM. One of the most important is the careful delineation of the roles of control processes and decision processes in retrieval. The application to part-list cuing has illustrated the importance of carefully specifying such factors. The sampling assumptions are another key element of the present approach in several different ways. First, the fact that sampling is probabilistic allows for a considerable degree of resampling in certain circumstances. Such resampling of previously sampled images is the basis for stopping the search and hence an important contributor to the limitations on retrieval. Second, the sampling equation (Equation 4) provides an explicit basis for combining cues and for focusing the search. The recovery equation (Equation 7) is also important, because it allows for imperfect recovery from sampled images but lets the amount of recovery depend on the cue-to-image strength. Thus the effects of such variables as presentation time are easily handled.

The use of a retrieval structure is an important part of the SAM theory. The retrieval structure abstracts from the LTS structure a simple enough representation that quantitative calculations are quite elementary and at the same time abstracts a complex enough representation that very elaborate situations can be handled sensibly by the theory. We admit that this approach hides many of the problems that arise during modeling of memory for highly structured materials by placing details of the interitem relationships inside the images, where they are dealt with during the recovery process. For complex structures, it might prove necessary to make the recovery process a very elaborate system. Nevertheless, the present approach does provide a vehicle by which such complex situations could be handled

within a search theory based on stochastic sampling, which we think is a significant step forward.

In comparison with theories that focus on the memory structure rather than the memory process, the SAM theory has several advantages. One is the ease with which forgetting phenomena may be handled. In this article, forgetting due to list length (or fanning) effects was predicted by SAMS. Forgetting due to delay is easily handled by modeling context changes. Interference phenomena can presumably be handled by a combination of these two factors, though such possibilities have yet to be explored in detail (but see Shiffrin, 1970, for one search model for interference effects). Another advantage is the emphasis on strategies that allow retrieval to be controlled by the subject.

Conversely, the greatest weakness of the SAM theory may be the freedom allowed by the numerous retrieval strategies. It would be tempting to predict any new result by arbitrarily assuming a new strategy designed to produce the observed effect. Fortunately, strategies may be manipulated experimentally, by instruction and task demand characteristics, so that models based on choices of strategies may be evaluated and tested.

Notwithstanding the success of SAMS in fitting free recall and part-list cuing, a general and complex model such as SAM cannot be judged a success simply on the basis of predicting the results from one or two studies or paradigms. We have therefore embarked on a program of evaluation in which the model (in essentially unchanged form) is applied to as many paradigms and results as possible. Reports on these applications are forthcoming. In particular, for applications to a wide variety of tasks involving free recall, categorized free recall, cued recall, and paired-associate recall, see Raaijmakers (1979) and Raaijmakers and Shiffrin (1980).

Reference Notes

1. Twohig, P. T. *Facilitation and inhibition of list words by list cues*. Paper presented at the meeting of the Eastern Psychological Association, Philadelphia, April 21, 1979.
2. Kintsch, W., & Kalk, J. M. Cueing with list items in multitrial free recall. In *Studies in mathematical learning theory and psycholinguistics* (CLIPR Pub-

lication No. 6). University of Colorado Computer Laboratory for Instruction in Psychological Research, April 1972.

References

- Allen, M. M. Cueing and retrieval in free recall. *Journal of Experimental Psychology*, 1969, 81, 29-35.
- Anderson, J. R. FRAN: A simulation model of free recall. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 5). New York: Academic Press, 1972.
- Anderson, J. R. *Language, memory, and thought*. Hillsdale, N.J.: Erlbaum, 1976.
- Anderson, J. R., & Bower, G. H. *Human associative memory*. Washington, D.C.: Winston, 1973.
- Atkinson, R. C., & Juola, J. F. Factors influencing speed and accuracy of word recognition. In S. Kornblum (Ed.), *Attention and performance IV*. New York: Academic Press, 1973.
- Atkinson, R. C., & Shiffrin, R. M. Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2). New York: Academic Press, 1968.
- Basden, D. R. Cued and uncued free recall of unrelated words following interpolated learning. *Journal of Experimental Psychology*, 1973, 98, 429-431.
- Basden, D. R., Basden, B. H., & Galloway, B. C. Inhibition with part-list cuing: Some tests of the item strength hypothesis. *Journal of Experimental Psychology: Human Learning and Memory*, 1977, 3, 100-108.
- Blake, M., & Okada, R. Intralist cuing following retroactive inhibition of well-learned items. *Journal of Experimental Psychology*, 1973, 101, 386-388.
- Bower, G. H. Stimulus-sampling theory of encoding variability. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory*. Washington, D.C.: Winston, 1972.
- Bower, G. H., Monteiro, K. P., & Gilligan, S. G. Emotional mood as a context for learning and recall. *Journal of Verbal Learning and Verbal Behavior*, 1978, 17, 573-585.
- Buschke, H. Short-term retention, learning, and retrieval from long-term memory. In D. Deutsch & J. A. Deutsch (Eds.), *Short-term memory*. New York: Academic Press, 1975.
- Crowder, R. G. *Principles of learning and memory*. Hillsdale, N.J.: Erlbaum, 1976.
- Estes, W. K. Statistical theory of spontaneous recovery and regression. *Psychological Review*, 1955, 62, 145-154.
- Glanzer, M., & Schwartz, A. Mnemonic structure in free recall: Differential effects on STS and LTS. *Journal of Verbal Learning and Verbal Behavior*, 1971, 10, 194-198.
- Hudson, R. L., & Austin, J. B. Effect of context and category name on the recall of categorized word lists. *Journal of Experimental Psychology*, 1970, 86, 43-47.
- Jones, G. V. A fragmentation hypothesis of memory: Cued recall of pictures and of sequential position. *Journal of Experimental Psychology: General*, 1976, 105, 277-293.
- Kintsch, W. Models for free recall and recognition. In

- D. A. Norman (Ed.), *Models of human memory*. New York: Academic Press, 1970.
- Kintsch, W. *The representation of meaning in memory*. Hillsdale, N.J.: Erlbaum, 1974.
- Luce, R. D. *Individual choice behavior: A theoretical analysis*. New York: Wiley, 1959.
- Luce, R. D., Bush, R. R., & Galanter, E. (Eds.), *Handbook of mathematical psychology* (Vol. 1). New York: Wiley, 1963.
- Mueller, C. W., & Watkins, M. J. Inhibition from part-set cuing: A cue-overload interpretation. *Journal of Verbal Learning and Verbal Behavior*, 1977, 16, 699-709.
- Murdock, B. B., Jr. The serial position effect in free recall. *Journal of Experimental Psychology*, 1962, 64, 482-488.
- Norman, D. A., & Rumelhart, D. E. *Explorations in cognition*. San Francisco: Freeman, 1975.
- Parker, R. E., & Warren, L. Partial category cuing: The accessibility of categories. *Journal of Experimental Psychology*, 1974, 102, 1123-1125.
- Raaijmakers, J. G. W. *Retrieval from long-term store: A general theory and mathematical models*. Doctoral dissertation, University of Nijmegen, Nijmegen, The Netherlands, 1979.
- Raaijmakers, J. G. W., & Shiffrin, R. M. SAM: A theory of probabilistic search of associative memory. In G. Bower (Ed.), *The psychology of learning and motivation*, Vol. 14. Academic Press, 1980.
- Ratcliff, R., & Murdock, B. B., Jr. Retrieval Processes in Recognition Memory. *Psychological Review*, 1976, 83, 190-214.
- Roberts, W. A. Free recall of word lists varying in length and rate of presentation: A test of total-time hypotheses. *Journal of Experimental Psychology*, 1972, 92, 365-372.
- Roediger, H. L., III. Inhibition in recall from cuing with recall targets. *Journal of Verbal Learning and Verbal Behavior*, 1973, 12, 644-657.
- Roediger, H. L., III. Inhibiting effects of recall. *Memory & Cognition*, 1974, 2, 261-269.
- Roediger, H. L., III. Recall as a self-limiting process. *Memory & Cognition*, 1978, 6, 54-63.
- Roediger, H. L., III, Stelson, C. C., & Tulving, E. Inhibition from part-list cues and rate of recall. *Journal of Experimental Psychology: Human Learning and Memory*, 1977, 3, 174-188.
- Roediger, H. L., III, & Thorpe, L. A. The role of recall time in producing hypermnnesia. *Memory & Cognition*, 1978, 6, 296-305.
- Rundus, D. Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior*, 1973, 12, 43-50.
- Shiffrin, R. M. Memory search. In D. A. Norman (Ed.), *Models of human memory*. New York: Academic Press, 1970.
- Shiffrin, R. M. Short-term store: The basis for a memory system. In F. Restle, R. M. Shiffrin, N. J. Castellan, H. Lindman, & D. B. Pisoni (Eds.), *Cognitive theory* (Vol. 1). Hillsdale, N.J.: Erlbaum, 1975.
- Shiffrin, R. M. Capacity limitations in information processing, attention and memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Vol. 4. Memory processes*. Hillsdale, N.J.: Erlbaum, 1976.
- Slamecka, N. J. An examination of trace storage in free recall. *Journal of Experimental Psychology*, 1968, 76, 504-513.
- Slamecka, N. J. Testing for associative storage in multitrial free recall. *Journal of Experimental Psychology*, 1969, 81, 557-560.
- Slamecka, N. J. The question of associative growth in the learning of categorized materials. *Journal of Verbal Learning and Verbal Behavior*, 1972, 11, 324-332.
- Smith, A. D. Output interference and organized recall from long-term memory. *Journal of Verbal Learning and Verbal Behavior*, 1971, 10, 400-408.
- Smith, E. E. Theories of semantic memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Vol. 6. Linguistic functions in cognitive theory*. Hillsdale, N.J.: Erlbaum, 1978.
- Smith, S. M. Remembering in and out of context. *Journal of Experimental Psychology*, 1979, 5, 460-471.
- Smith, S. M., Glenberg, A. M., & Bjork, R. A. Environmental context and human memory. *Memory & Cognition*, 1978, 6, 342-353.
- Tulving, E. Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press, 1972.
- Tulving, E. Cue-dependent forgetting. *American Scientist*, 1974, 62, 74-82.
- Tulving, E., & Osler, S. Effectiveness of retrieved cues in memory for words. *Journal of Experimental Psychology*, 1968, 77, 593-601.
- Tulving, E., & Pearlstone, Z. Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 1966, 5, 381-391.
- Watkins, M. J. Inhibition in recall with extralist "cues." *Journal of Verbal Learning and Verbal Behavior*, 1975, 14, 294-303.
- Watkins, M. J., & Watkins, O. C. Cue-overload theory and the method of interpolated attributes. *Bulletin of the Psychonomic Society*, 1976, 7, 289-291.
- Waugh, N. C. Presentation time and free recall. *Journal of Experimental Psychology*, 1967, 73, 39-44.
- Waugh, N. C., & Norman, D. A. Primary memory. *Psychological Review*, 1965, 72, 89-104.
- Williams, M. D. *Four observations on the process of retrieval from very long-term memory*. (Doctoral dissertation, University of California, San Diego, 1977). *Dissertation Abstracts International*, 1978, 38(9-B), 4517.
- Wood, G. Retrieval cues and the accessibility of higher order memory units in free recall. *Journal of Verbal Learning and Verbal Behavior*, 1969, 8, 782-789.

Received September 27, 1979

Revision received April 20, 1980 ■