Psychological Review

VOLUME 84 NUMBER 2 MARCH 1977

Controlled and Automatic Human Information Processing: II. Perceptual Learning, Automatic Attending, and a General Theory

Richard M. Shiffrin
Indiana University

Walter Schneider
University of California, Berkeley

The two-process theory of detection, search, and attention presented by Schneider and Shiffrin is tested and extended in a series of experiments. The studies demonstrate the qualitative difference between two modes of information processing: automatic detection and controlled search. They trace the course of the learning of automatic detection, of categories, and of automatic-attention responses. They show the dependence of automatic detection on attending responses and demonstrate how such responses interrupt controlled processing and interfere with the focusing of attention. The learning of categories is shown to improve controlled search performance. A general framework for human information processing is proposed; the framework emphasizes the roles of automatic and controlled processing. The theory is compared to and contrasted with extant models of search and attention.

I. Introduction

In Part I of this paper (Schneider & Shiffrin, 1977) we reported the results of several experiments on search and attention that led us to formulate a theory of information processing based on two fundamental processing modes: *controlled* and *automatic*. In the context of search studies, these modes took the form of *controlled search* and *automatic detection*. Controlled search is highly demanding of attentional capacity, is usually serial in nature with a limited comparison rate, is easily established, altered, and even reversed by the subject, and is strongly dependent on load. Automatic detection is relatively well learned in long-term memory, is demanding of attention only when a target is presented, is parallel in nature, is difficult to alter, to ignore, or to suppress once learned, and is virtually unaffected by load.

In the present article we shall report several studies to elucidate further the properties of automatic and controlled processing and their interrelations, to demonstrate the qualitative difference between these processing modes, to study the development of automatic detection and the role of the type and nature of practice in such development, to study the effects of categorization, and to examine the development of automatic attending and its effects. After the presentation of the studies we shall present a general theory of information processing, with emphasis on the

The research and theory reported here were supported by PHS Grant 12717 and a Guggenheim Fellowship to the first author, Grant MH23878 to the Rockefeller University, and a Miller Fellowship at University of California at Berkeley to the second author. This report represents equal and shared contributions of both authors.

Requests for reprints should be sent to Richard M. Shiffrin, Department of Psychology, Indiana University, Bloomington, Indiana 47401.

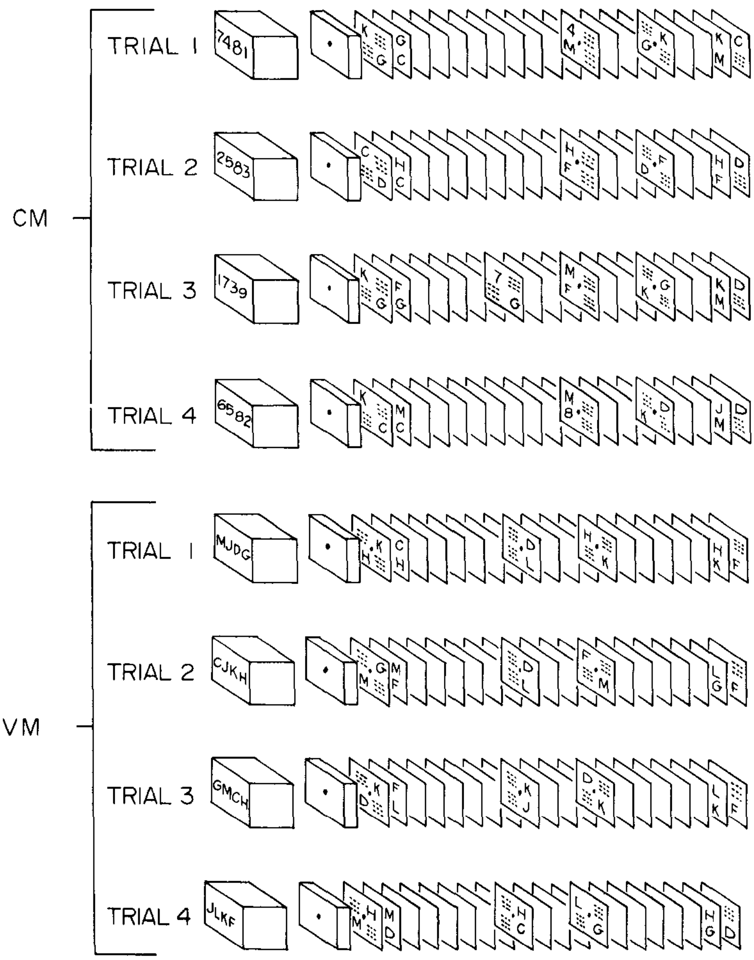

Figure 1. Examples of trials in the multiple-frame search paradigm of Experiment 1, Part I. In all cases, memory-set size = 4 and frame size = 2. Four varied-mapping (VM) trials and four consistent-mapping (CM) trials are depicted. The memory set is presented in advance of each trial, then the fixation dot goes on for .5 sec when the subject starts the trial, and then 20 frames are presented at a fixed time per frame. Either 0 or 1 member of the memory set is presented during each trial. Frame time, memory-set size, and frame size are varied across conditions.

roles of automatic and controlled processing. Our theory will then be compared and contrasted with extant theories of search and attention.

A. *Review of Paradigms and Results From Part I*

The paradigm for most of the studies of Part I (and Part II also) is depicted in Figure 1. Four elements are presented simultaneously in a square; and their joint presentation for a brief period of time is termed a *frame*. Each trial consists of the presentation of 20 frames in immediate succession. The elements presented are characters (i.e., digits or consonants) or random dot masks. In advance of each trial the subject is presented with several items, called the *memory set*, and is then required to detect any memory-set items that appear in the subsequent frames. The frame time is kept constant across the 20 frames of each trial, and the basic dependent variable is the psychometric function relating accuracy to frame time for each condition.

Three basic independent variables were manipulated in Part I/Experiment 1. The number of characters in each frame (the frame size, F) was varied from 1 to 4 (but was constant for all frames of a given trial). The number of characters presented in advance of a trial (the memory-set size, M) varied from 1 to 4. The product of M and F is termed the *load*. A memory-set item that appears in a frame is called a *target;* an item in a frame that is not in the memory set is called a *distractor*. One half of the trials contained one target, and one half contained no target. Finally, and most important, the nature of the training procedure across trials and the relation of the memory-set items to the distractors were varied. In the *consistent mapping* (CM) procedure, across all trials, memory-set items were never distractors (and vice versa). In addition, memory-set items were from one category (e.g., digits) and distractors from another category (e.g., consonants). In the *varied mapping* (VM) procedure, memory-set items and distractors were randomly intermixed over trials and were all from one category (e.g., consonants).

Figure 1 gives examples of trials in both the VM and CM conditions. Depicted are four consecutive trials from a CM block and four from a VM block in each of which $M = 4$ and $F = 2$. Table 1 gives the memory set, distractor, and target (if present) for each trial in Figure 1. Note that memory-set items and distractors intermix across trials in the VM condition but do not intermix over trials in the CM condition.

The most important results are shown in Figure 2. These results showed that the VM conditions were strongly affected by load and were quite difficult; the CM conditions were virtually unaffected by load and were all easier than even the easiest VM condition. It was suggested that a controlled, serial search was operating in the VM conditions and that a qualitatively different process, automatic detection, was operating in the CM conditions.

The results of Part I/Experiment 1 were analyzed on the basis of the accuracy of the detection responses. Part I/Experiment 2 utilized comparable conditions but presented only a single frame on each trial; accuracy was high and the results were analyzed on the basis of the reaction time of the responses. The results confirmed those of Part I/Experiment 1, and a quantitative model was fit to the VM results of both experiments. This model assumed that controlled search was a serial, terminating comparison process in which one first compared all frame items against one memory-set item before switching to the next memory-set item. Each comparison and each switch was assumed to require some time to be executed. The success of the model in fitting the results of both experiments suggests that the same search mechanisms underlie search experiments that utilize both accuracy and reaction time measures, and suggests that the same search mechanisms underlie performance in both divided-attention and search paradigms.

The vast differences between results of the CM and VM conditions provided a basis for reorganizing and classifying the results of previous search and detection studies. These studies fell into a relatively simple organization, and many perplexing and seemingly contradictory results became explicable.

Part I/Experiment 3 utilized a multiple-target, multiple-frame procedure. The study was similar to Part I/Experiment 1, but the subject was presented either zero, one, or two targets per trial and was required to report the number of detected targets. The condi-

Table 1
Examples of CM and VM Trials for Four Successive Trials

Trial	Memory set	Distractor set	Target
	Consistent mapping (CM)		
1	7481	KGJCM	4
2	2583	CHFLD	none
3	1739	KGFDM	7
4	6582	CMJKD	8
	Varied mapping (VM)		
1	MJDG	CFHKL	D
2	CJKH	LGDFM	none
3	GMCH	DLFKJ	none
4	JLKF	CDGHM	L

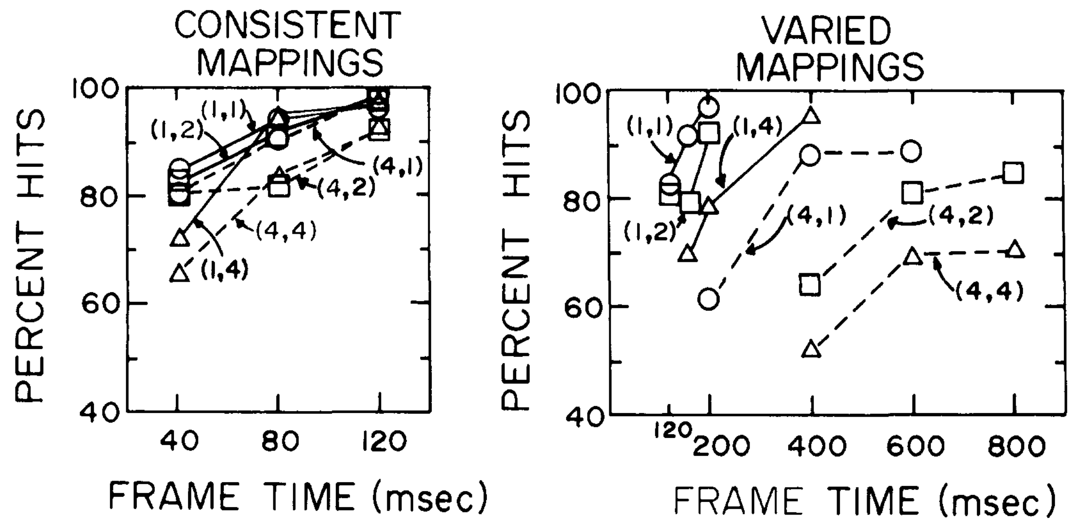

Figure 2. Data from Experiment 1, Part I. Hits as a function of frame time for each of the 12 conditions. Three frame times were utilized for each condition. The first number in parentheses indicates the memory-set size and the second number in parenthesis indicates the frame size.

tions of greatest interest were those in which two targets per trial were presented. In these cases the spacing between the two targets varied from 0 to 4 (0 spacing indicates simultaneous presentation; spacings of 1, 2, and 4 indicate the number of intervening frame plus 1). Target similarity also varied, the two targets being either physically identical (II) or different (NI).

The results for Part I/Experiment 3 showed markedly different patterns for the VM and CM conditions. Consider target similarity first. In the VM conditions detection of identical targets (II) was superior to detection of different targets (NI). However, in the CM conditions, the reverse was true: NI detection was superior to II detection. Consider target spacing next. In the VM conditions performance was lowest when the targets occurred in successive frames (spacing 1). However, in the CM conditions performance was lowest when the targets were simultaneous (spacing 0). These results helped emphasize the qualitative difference between the automatic detection processing mode presumed to be utilized in the CM conditions, and the controlled search mode presumed to be utilized in the VM conditions. In certain of the studies in this paper, we shall utilize this multiple-target, multiple-frame procedure and will infer the presence of automatic detection or controlled search from the nature of the spacing effect and the target-similarity effect.

B. *Rationale for the Experiments*

We argued in Part I that the cause of the difference between the CM and VM results was the consistency of the mapping (over trials) of the memory-set items and distractors to responses. We argued that consistent mapping leads to the development of automatic detection, which enables automatic-attention responses to become attached to the memory-set items. The automatic-attention responses enable the serial, controlled search to be bypassed by a parallel detection process unaffected by load. However, these hypotheses must be further tested for the following reasons. First, the fact that memory-set items were categorically distinct from the distractors was confounded with the consistency of the mapping. These factors will be separated in Part II/Experiments 1, 2, and 3. Second, the course of development of the hypothesized automatic detection process was not studied in Part I. The course of learning will be traced in Part II/Experiments 1 and 3. Third, there was no demonstration in Part I that an automatic-attention response is learned in CM paradigms. Such a fact will be suggested by Part II/Experiments 1 to 3 and demonstrated in Part II/Experiment 4. Finally, although these goals provide an initial justification for the present experiments, these studies will serve an even more important purpose in elucidating the characteristics and development of automatic processing and in differentiating automatic from controlled processing.

II. The Development of Automatic Processing: Perceptual Learning

A. *Perceptual Learning and Unlearning in a CM Task Using Letters Only*

Experiment 1 of the present series is simple in conception. With the use of the same basic multiple-frame paradigm used in Part I/Experiment 1 (see Figure 1), performance is examined as a function of the amount of training under consistent-mapping conditions. One change is made, however: Both the distractor set and the memory ensemble consist of consonants. This procedure enables us to study the acquisition of automaticity when the two sets are not already-learned categories. Furthermore, it had been observed informally that the use of digits and consonants (in the CM conditions of Part I/Experiment 1) led to a relatively rapid acquisition of automatic detection. It was felt that the use of letters only to make up the two sets would slow down acquisition.

Values of M and F were chosen so as to make controlled processing difficult ($M = 4$, $F = 2$), and a frame time (f) of 200 msec was chosen so as to make performance low when controlled processing was being utilized. (These choices are justified by the data in Figure 2.) Thus it was expected that performance would be quite poor at the start of training, when automatic detection had not been learned and controlled search had to be used, but would improve markedly as automatic detection developed.

1. *Method*

The CM presentation procedure of Part I/Experiment 1 was utilized (see Figure 1). The frame size was always equal to 2 and the memory-set size was always equal to 4. Two disjoint character sets were used for the memory ensemble and the distractor set. One consisted of the following nine consonants: B, C, D, F, G, H, J, K, and L; the other consisted of the following nine consonants: Q, R, S, T, V, W, X, Y, and Z.

There were four new subjects, two of each sex, all naive to our tasks.[1] Two subjects began training with the consonants from the second half of the alphabet as the memory ensemble, and two subjects began with the consonants from the first half of the alphabet as the memory ensemble. There were five blocks of trials per session, each containing 60 test trials. There were no practice trials. Subjects were informed at the start of the experiment concerning the nature of the study and the composition of the memory ensemble and the distractor set. Each trial began with the presentation of a memory set ($M = 4$) selected randomly from the memory ensemble.

During the first 1,500 trials of the experiment for each subject, the frame time was 200 msec; during the following 600 trials, 120 msec. After these 2,100 trials the memory ensemble and the distractor set were switched for each subject and the frame time was set back to 200 msec. This reversal condition was then run at the frame time of 200 msec for a total of 2,400 additional trials. The subjects were informed of the reversal at the time of the switch.

The subjects responded whenever a target was detected, or gave a negative response at the end of the trial. The accuracy of the response and the reaction time for hits were both recorded. Subjects heard a tone signifying an error after each incorrect response.

2. *Results and Discussion*

The results are presented in Figure 3. Results for each block, averaged across subjects, are graphed consecutively. Thus, the graphed points in each interval are based on 120 observations.

In the initial group of 1,500 trials, the hit rate rose from just over 50% to about 90%, while the false alarm rate dropped from about 12% to about 3% (and the reaction

---

[1] Since the experiments are not reported in their chronological order, we list here the experiments in their original order and indicate the subjects that were used in each. Experiment 1 was run with new subjects. Experiment 4 was run next, using the subjects who took part in Experiment 3 of Part I. Experiment 2 was then run on three of the four subjects in Experiment 4. Experiment 3 was run last, on new subjects.

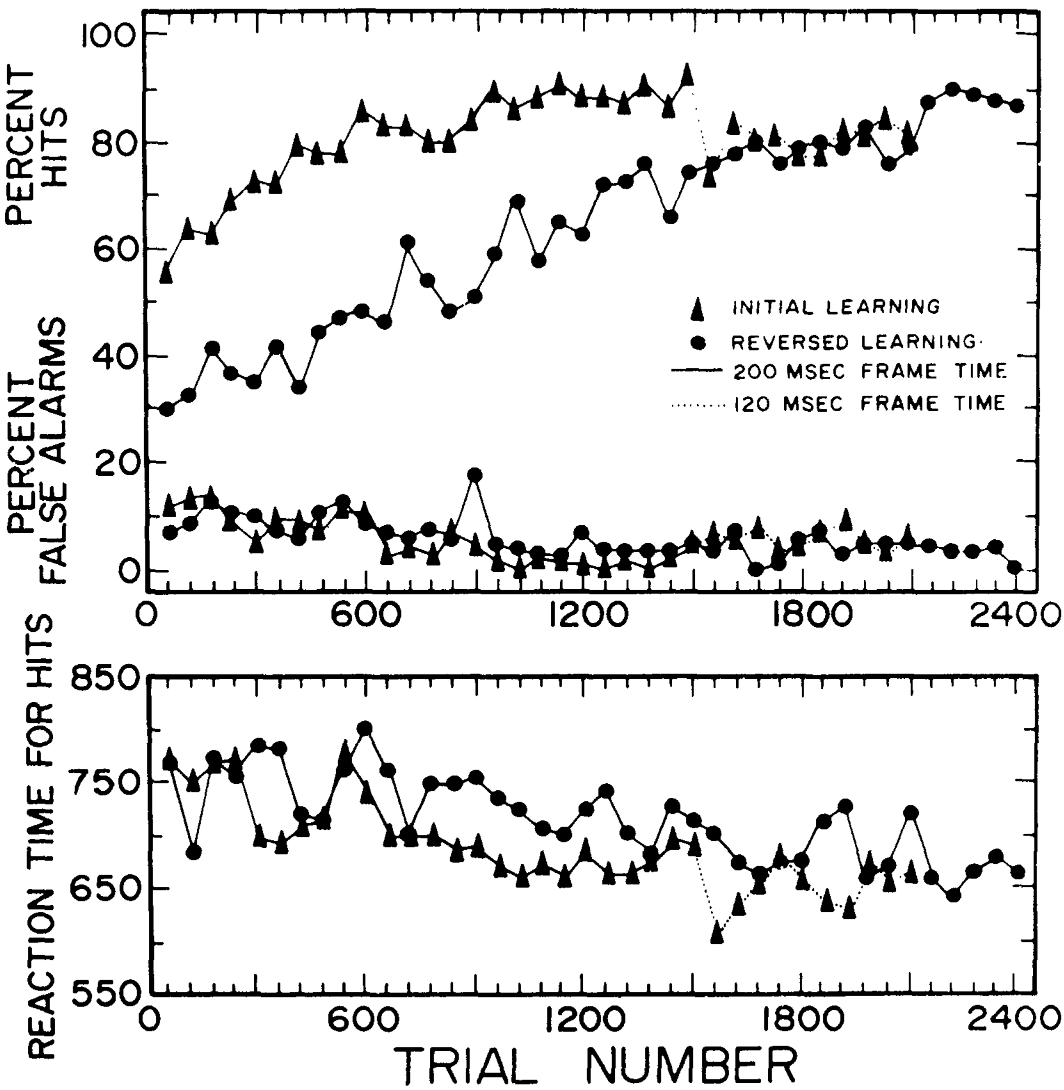

Figure 3. Data from Experiment 1. Initial consistent-mapping learning and reversed consistent-mapping learning for target and distractor sets taken from the first and second halves of the alphabet. Memory-set size = 4, frame size = 2, frame times are shown. Percentage of hits, percentage of false alarms, and mean reaction time for hits are graphed as a function of trial number. After 2,100 trials the target and distractor sets were switched with each other.

time for hits dropped from about 770 msec to about 670 msec).

It seems plausible that the subjects adopted a controlled search strategy at the start of training. Some support for this hypothesis is found in a comparison of the data from this study with the data from Experiment 3 of Part I (though different subjects were involved in the two studies). The hit and false alarm probabilities in the first 60 trials of the present study were .56 and .13, respectively. The estimated probability of detecting a single target and the observed probability of giving a false alarm when no target was present were .60 and .18, respectively, in the VM condition of Experiment 3b of Part I in which M was equal to 4 and F was equal to 2 (this condition and the first 60 trials of the present study were both run at a frame time of 200 msec). These two conditions should give similar results if controlled search is being utilized in both studies; in fact, the probabilities are quite similar.

It should be noted that the subjects all reported extensive, attention-demanding rehearsal of the memory set during about the first 600 trials of Experiment 1, but they gradually became unaware of rehearsal or other attention-demanding controlled processing after this point. Both the objective evidence and subjective reports, then, suggest that the subjects used controlled search at the start of Experiment 1 but gradually shifted to automatic detection.

The hit rate and false alarm rate after 1,500 trials appeared to be changing only slowly, but this could have resulted from a ceiling effect. The frame time was therefore reduced to 120 msec for the next 600 trials. At the conclusion of this additional training, the hit rate appears to have reached about 82%, while the false alarm rate dropped to about 5%. These results may be compared with the CM results from Experiment 1 of Part I (see Figure 2). In that study, with $f = 120$ msec, $M = 4$, and $F = 2$, the CM hit rate was about 92% and the false alarm rate (not shown) was about 5%.

Thus, the present results indicate that subjects were utilizing automatic detection by the end of the initial 2,100 trials of CM training. First, they performed at a level not far from that of the Part I/Experiment 1 subjects in the comparable CM condition with the same frame time. Furthermore, the present study utilized a larger memory ensemble and distractor set than the VM conditions of Part I/Experiment 1, a fact that could only have increased the task difficulty. Second, from the Part I/Experiment 1 VM data we can estimate that a frame time of 600–800 msec would have been required to achieve a similar level of performance had controlled search been used.

At this stage of the experiment, after 2,100 trials of CM training, the subjects could be expected to have developed a well-learned attention response to the members of the memory ensemble. If such learning is firmly planted in long-term memory and if the attention response occurs automatically, then it should prove very difficult for the subject to alter or unlearn his automatic response in any short period of time. To test this hypothesis we reversed the memory ensemble and the distractor set for each subject (and set the frame time back to 200 msec). We hypothesized that automatic detection would prove impossible and that the subject would be forced to revert to controlled search.

The results of the reversal were quite dramatic. The hit rate just after reversal fell to a level well below that seen at the start of training when the subjects were completely unpracticed. Very gradually thereafter the hit rate recovered, so that after 2,400 trials of reversal training, performance reached a level about equal to that seen after 1,500 trials of original training. In summary, the original training resulted in quite strong negative transfer, rather than positive transfer.

Subject's verbal reports indicated that a shift back to controlled search occurred after reversal. Initially, before reversal, all subjects reported rehearsing the memory set during each trial. After the 2nd day (600 trials) subjects reported that they were no longer rehearsing and only glanced at the memory set. However, subjects reported that after reversal they tried various methods to perform the task and eventually ended up rehearsing the memory-set items again, though this rehearsal also decreased after a week of postreversal practice. This pattern of reports could indicate a shift from controlled search to automatic detection during original learning, then a shift after reversal to controlled search, and then finally a return to automatic detection.

What might be the cause of the negative transfer after reversal? If the reversal caused subjects to revert to controlled search, then it might have been expected that performance would fall to the level seen at the start of original learning (when controlled search was presumably utilized). The actual results suggest that controlled search is hindered when the distractors are items that subjects have been previously trained to respond to as

CM targets (i.e., items that, due to previous training, give rise to automatic-attention responses). This hypothesis is verified in Experiment 4, to be reported later.[2]

Beyond the poor performance at the start of reversal learning, negative transfer is also evidenced by the extremely large amount of training needed to overcome the effects of original learning. After about 900 trials of original practice the hit rate reached about 90%. However, it took about 2,100 trials to return to this performance level after reversal. Apparently the necessity to unlearn the attention responses to the previous memory-set items, or the necessity to overcome some learned inhibition to the previous distractors, or both, causes a reduction in the rate of acquisition of automatic processing after reversal.

It would be interesting to know whether negative transfer would obtain in several variations of our paradigm. For example, in the present study, the initial training might have caused the memory set to become a learned category, but probably did not cause categorization of the distractor set. If the distractor set was a well-known category, would reversal still cause a severe performance drop? Furthermore, it could be argued that the initial training did not continue long enough for the negative transfer to reach full strength. The amount of negative transfer might depend on the amount of overtraining during initial learning. Such a view has been put forth for certain verbal and animal learning situations (see Mandler, 1962, 1965; Jung, 1965 for two views on the subject). Thus it might be asked whether performance drops would be caused by reversal if the original automatic detection responses were greatly overlearned. These two questions are answered by Experiment 3.

B. *The Reversal of Consonants and Digits for Well-Practiced Subjects*

Throughout the series of experiments in Part I, the memory ensemble in the CM condition was always the same, either digits or consonants, and no member of the CM memory ensemble was ever a distractor, even in the VM conditions. The VM conditions always consisted solely of items from the distractor set in the CM conditions. Thus a subject at the conclusion of the series of Part I studies who had been searching, say, for consonants in digit distractors, had never been given a consonant as a distractor: Whenever a consonant had appeared, it was a target. These facts naturally suggested a study in which the roles of consonants and digits were reversed for these subjects.

1. *Method*

The procedure was identical to the multiple-frame, multiple-target paradigm of Experiment 3a of Part I (whose results were described in the Introduction) in most respects, except that the memory ensemble and the distractor set were switched in both the CM and VM conditions. Thus in the CM condition the two subjects who had been searching for consonants in digit distractors now searched for digits in consonant distractors; in the VM condition these two subjects had been searching for digits in digits and now searched for consonants in consonants. These contingencies were reversed for the remaining two subjects who had previously been searching for digits in consonant distractors in their CM conditions.

The frame time (f) was 60 msec for the CM conditions and 200 msec for the VM conditions. Just as in Experiment 3a of Part I, M was equal to 2, and F was equal to 2; zero targets were presented on one fourth of the trials, one target was presented on one fourth of the trials and two targets were presented on one half of the trials. When two targets were presented, they were either identical (II) or nonidentical (NI), and the spacing between them was either 0, 1, 2, or 4. Each subject was given 12 blocks of 132 trials in each of the VM and CM conditions. The VM and CM blocks alternated.

Upon completion of the initial set of 24 blocks, the frame time was increased to 120 msec and an additional 18 blocks of CM trials were run for each subject.

[2] One might hypothesize that the basic comparison rate of controlled search is altered and slowed after reversal, thereby accounting for the negative transfer. Two facts argue against this hypothesis. Experiment 2 shows that controlled search proceeds at an unchanged rate even when all items consist of targets from previous CM conditions. Experiment 4 shows that even one CM target presented in a to-be-ignored spatial location greatly reduces detection of a simultaneously presented target in a to-be-attended location. This result suggests that reversal performance is harmed because attention tends to be drawn to the distractors.

It should be noted that the subjects for this experiment had had roughly 20,000 trials of practice in the various CM and VM conditions in earlier studies (see Footnote 1).

2. *Results and Discussion*

The results are depicted in Figure 4. A simple correction for guesses and false alarms has been carried out on the raw data, as described in Part I, but this correction did not affect the qualitative features of the results. The circles in the left-hand panel give the comparable VM prereversal data from Experiment 3a of Part I ($f = 200$ msec), while the circles in the center panel give the comparable CM prereversal data from Experiment 3c of Part I ($f = 60$ msec); the triangles in the two left-most panels give the data from the present experiment (see Footnote 1).

The most dramatic findings of Experiment 2 are shown in the middle panel of Figure 4, which gives the CM reversal results. It is evident that performance drops off markedly after reversal. Estimated percentage of detection of two targets is about 80% prior to reversal and about 20% after reversal. Furthermore, these data represent 12 blocks of training over which very little, if any, recovery was taking place. It is interesting to note that all of the subjects predicted that no change in their performance would occur after the consonants and digits were reversed, and the subjects were uniformly startled and even dismayed by the extreme difficulty of the reversed task.

These results resolve the confounding between categorization and the mapping conditions. Even when there is a well-learned and well-practiced categorical difference between memory ensemble and distractor set, reversal in the CM paradigm gives rise to marked impairment of performance. Thus, categorical differences between memory ensemble and distractor set cannot be the key factor underlying the utilization of automatic detection. Rather it is the consistency of mapping that underlies the acquisition of automatic detection.

Consider next the VM data given in the

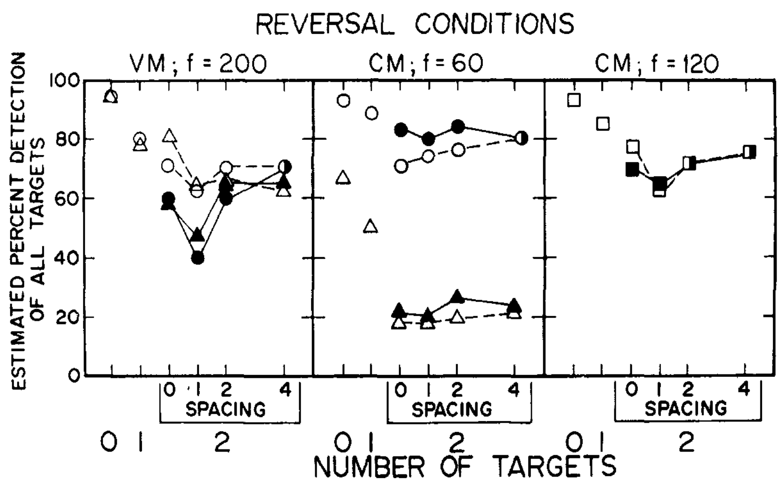

Figure 4. Data from Experiment 2. Probabilities of correctly detecting one target when one was presented, or two targets when two were presented, as a function of spacing, for each condition. The data shown have been adjusted to remove effects of guessing and false alarms. The observed percentage of "none" responses to no-target trials is also shown. Circles show the estimated percentage of detection prior to reversal (data from Experiment 3, Part I), while triangles indicate performance after reversal in Experiment 2. The squares in the right-hand panel show post-reversal detection from Experiment 2 after frame time was increased to 120 msec. The open symbols and dashed lines indicate the II conditions (identical targets); the closed circles and solid lines indicate the NI conditions (different targets). (VM = varied mapping; CM = consistent mapping; f = frame time, in msec.)

left-hand panel of Figure 4. It is evident that the switch from consonants to digits (or vice versa for other subjects) had almost no effect. The levels of performance remain the same as in Experiment 3a of Part I, and the qualitative pattern of results is still about the same, showing the pattern indicative of controlled search. That is, worst performance occurred at spacing 1, and performance was worst when the targets were different (NI). These effects indicate the presence of controlled search. In summary, a change of the character set used in the VM conditions had no appreciable effect under the present conditions.

These results are most interesting, because after reversal the memory ensemble and the distractor set both consisted of items that had come to elicit automatic-attention responses in previous training. In the reversal conditions of Experiment 1 we saw that making the distractor set but not the memory ensemble consist of such "automatic" items resulted in considerable negative transfer. In the present experiment it is evident that performance is not worse after reversal when subjects have been trained to respond to *both* the targets and distractors in previously CM training.

These results might be expected on the basis of the following reasoning. Just after reversal in both Experiment 1 and in the CM conditions of Experiment 2 the subjects were using controlled search. In Experiment 1, after reversal, the automatic-attention responses to the distractor items tended to draw attention away from target items to which subjects had not previously been trained to respond and hence to order the controlled search in such a way that the distractor items were compared earlier in the search. Thus performance was impaired. (This hypothesis is confirmed in Experiment 4.) However, in the present experiment, after reversal, all items, distractors and targets alike, gave rise to attention responses, which could be expected to cancel each other out on the average. That is, there should be no selective bias to compare distractors prior to targets since both types of items give rise to equivalent attention responses. Thus, the controlled search operates normally, with targets compared in a random order of search, and performance is equivalent to that seen prior to the switch of character sets.

In short, the use in a search task of stimuli that elicit automatic-attention responses (developed through previous CM training) can (a) improve performance over that seen in the usual VM conditions if these stimuli are memory-set items but not distractors, in which case automatic detection will be utilized; (b) leave performance unchanged from that seen in the usual VM conditions if these stimuli are both memory-set items and distractor items, in which case normal controlled search will be utilized; and (c) impair performance from that seen in the usual VM conditions if these stimuli are distractors but not memory-set items, in which case subjects will utilize controlled search that will be reduced in effectiveness by automatic attending to distractors.

Let us now turn to the results of the final reversed CM training blocks. We suspected that the subjects in the reversed CM condition (the middle panel) were reverting back to the use of controlled search, but performance levels were so low that it was difficult to interpret the pattern of results. Thus frame time in the CM reversal condition was increased to 120 msec, and 18 blocks of additional CM training were run.

The results of the trials with $f = 120$ msec are represented by the squares in the right-hand panel of Figure 4. It should be noted first that the pattern of the results is indicative of controlled search in that the worst performance occurs at spacing 1. However, the results do not show the target-similarity effect that we have been led to expect (in Experiment 3 of Part I) for controlled search: Results for conditions using II target pairs are not superior to those using NI target pairs. Furthermore, the overall level of performance is higher than one would expect for the VM conditions. At a frame time of 120 msec performance is higher than that seen for the VM conditions at a frame time of 200 msec (in the left-hand panel of Figure 4).

The results can be explained simply and

elegantly by the hypothesis that the subjects were carrying out a controlled search for the category as a whole. The Experiment 2 reversal condition differed from the Part I/Experiment 3 VM conditions in that the memory-set items and the distractor items in the present instance were categorically different (numbers and letters). This categorical difference obviously did not allow the subjects to utilize automatic detection after reversal. However, the categorical difference could very well have helped the subjects to adopt a more efficient controlled search. Suppose the subject ignored the specific members of the memory set and searched instead for any instance of the category "numbers" (or "letters" as the case may be). Then only one comparison would be needed for each display item in each frame, regardless of memory-set size.

Compared with the case when a categorization is unavailable, a controlled search for the category of each input saves search time in two ways. First, there is no need for multiple memory-set comparisons for each frame item. Second, there is no need for time-consuming switches between memory-set items. Thus performance improves considerably. Furthermore, the distinction between identical and nonidentical multiple targets is no longer meaningful—both targets are simply category members, regardless of their physical similarity. Thus no difference between the II and NI conditions would be predicted. The results in Figure 4 support these contentions.

As far as we can tell, through a block-by-block analysis of the reversal results, the subjects did not begin to recover any appreciable degree of automatic detection even after 30 blocks of CM reversal training. Undoubtedly, the difficulty in relearning (or unlearning) is related to the great amount of overtraining with the original mapping of stimuli to responses. Eventually, however, were the reversal experiment continued, we would expect automatic detection to develop once again.

In summary, then, the results of Experiment 2 show that a switch of character sets does not affect VM performance, but a switch of memory ensemble and distractor set causes a great decrement in performance in the CM conditions. The CM decrement occurs despite the categorical difference between the memory-set items and distractor items (letters vs. numbers). It is argued that the subjects, after reversal, adopt a controlled search strategy in which they search for the category as a whole rather than search for the individual members of the category.

The results of Experiments 1 and 2 could hardly have provided a more dramatic demonstration of the qualitative difference between automatic detection and controlled search, and of the dependence of the search mechanism on the nature of the training procedures (i.e., the consistency of the stimulus–response mapping).

The results of Experiments 1 and 2 also demonstrate the long-term nature of the learning underlying the automatic process. Experiment 1 showed that even original learning of an automatic response could take thousands of trials of practice to develop if a previously known categorization did not distinguish the memory ensemble and the distractor set. Relearning after reversal was even slower to develop. Experiment 2 showed that relearning after reversal could be very retarded even when the sets were categorized, as long as the original automatic detection process was well overlearned.

C. *The Role of Categorization in Controlled Search and in the Development of Automaticity*

The role played by categorization in controlled search and in the development of automatic detection is suggested by results of the preceding experiments, but it is an important enough concept to be examined in a separate experiment. Furthermore, the nature of the process by which categories are learned is also worth studying. These considerations underlay the design of Experiment 3. We decided to use new subjects and to train them in two conditions, in neither of which were the memory and distractor sets preexperimentally categorized. The first condition was designed to induce controlled

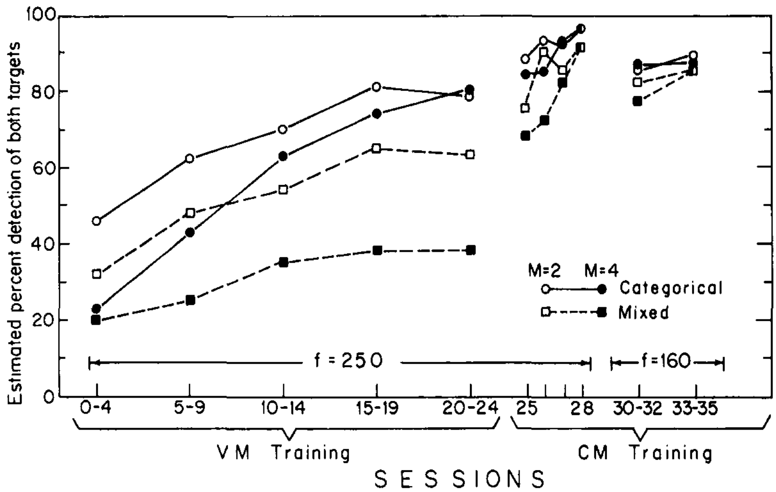

Figure 5. Data from Experiment 3. Estimated percentage of detection of both targets, as a function of session number, averaged over spacing and target-similarity conditions, for each of the four main conditions. At Session 25 training was switched to a CM procedure. At Session 30, the frame time was reduced to 160 msec. (VM = varied mapping; CM = consistent mapping; M = memory-set size; f = frame time, in msec).

search in circumstances such that a categorization could not develop. The second condition was also designed to induce controlled search, but in circumstances such that a categorization of the memory sets could be learned. Both of these conditions utilized a varied-mapping procedure to prevent automatic detection, and both were followed by a series of CM trials, during which trials automatic detection developed. Thus, the CM trials traced the course of development of automatic detection when a categorization was, or was not, present at the start of CM training.

1. *Method*

The subjects each took part in two conditions called *mixed* and *categorical*, run in different blocks. In each condition the memory ensemble and distractor set were drawn from a total of eight consonants. The two eight-consonant sets did not overlap. The sets were {GMFP, CNHD} and {RVJZ, BWTX}. The assignment of these two sets to the two conditions was permuted across subjects.

Both conditions used a VM procedure similar in certain respects to that of Experiment 2 of this report and Experiment 3 of Part I. The multiple-target, multiple-frame paradigm was used with $f = 250$ msec per frame. Only 12 frames were used on each trial, with all targets appearing on frames 3 through 10. Frame size was equal to 2 in all conditions, and two memory-set sizes were used in different blocks: $M = 2$ and $M = 4$.

In the mixed condition, the members of the memory set were chosen randomly *on each trial* from the ensemble of eight consonants; the distractor set consisted of four items randomly chosen from those consonants not used in the memory set. These four distractors were chosen randomly to fill the nontarget, nonmask, positions in the various frames of the trial. The key feature of the mixed condition was the fact that memory-set items and distractors were randomly intermixed from trial to trial.

In the categorical condition, the eight ensemble items were divided into two sets of four that remained disjoint throughout the experiment for each subject. The sets of four are indicated by the position of the comma in the above listing of the consonants making up each set. Note that the two sets of four were chosen to be highly confusable with each other in the sense that pairs of visually confusable letters were separated into the two sets. On a given trial one of these sets of four was chosen randomly to be the distractor set, and the memory-set items were then chosen randomly from the remaining set. Thus, in the categorical condition there were just two disjoint categories of items, one from which the memory set was drawn, and one that served as the distractor set. However, the category that served as the distractor set varied randomly from trial to trial. Thus, the procedure still involved a varied mapping, but the categories were never intermixed and could eventually become well learned.

To make it easier for the subjects to learn the categories, the members of the memory set presented at the start of each trial were always presented in alphabetical order (this was also done in the mixed condition). In each session there were four blocks, always run in the following order: (a) categorical condition, $M = 2$; (b) categorical condition, $M = 4$; (c) mixed condition, $M = 2$; (d) mixed condition, $M = 4$. There were 96 trials in each block, 24 with no target, 24 with one target, and 6 for each of the eight combinations of spacing and target similarity when two targets were presented.

Four new subjects, one male and three females, took part. Each ran in 24 sessions of about 1 hour each. This finished the first, varied mapping, phase of the study. After these 24 sessions the subjects should have been using controlled search in both conditions, but should have learned the categories in the categorical condition. Before turning to the CM procedure used in subsequent sessions, we shall discuss the VM results.

2. *Results and Discussion of the VM Conditions*

The results are summarized in Figures 5 and 6. Figure 5 gives the estimated probability of detecting both targets when two are present, averaged across target similarity, spacing, subjects and certain combinations of sessions. Performance in all four main conditions tended to improve during VM training, but is clearly stable over the last 10 sessions (15–24). Initially the $M = 2$ condition is superior to the $M = 4$ condition in both the categorical and mixed cases. After training, however, the pattern changes considerably. In the categorical conditions

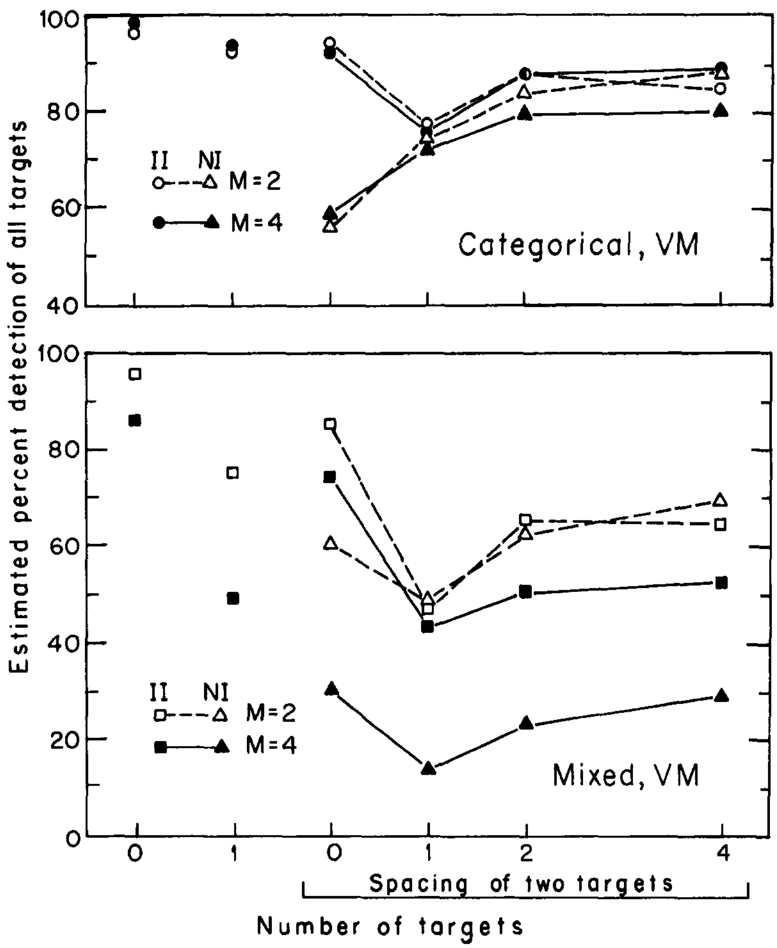

Figure 6. Data from Experiment 3. Estimated percentage of detection as a function of number, spacing, and similarity of targets for VM Sessions 19–23 for all conditions. (VM = varied mapping; M = memory-set size; II = identical targets; NI = nonidentical targets.)

the results of the $M = 2$ and $M = 4$ conditions converge (indicating category learning and the use of a categorical search strategy). In the mixed condition performance in the $M = 4$ condition remains much lower than in the $M = 2$ condition, but performance in both of these is worse than in either categorical condition (suggesting that the presence of well-known categories reduces the effective memory-set size to 1).

To ascertain whether controlled search was being utilized in these various VM conditions and to determine the nature of that search, it is necessary to consider the spacing functions shown in Figure 6. This figure gives the spacing functions for Sessions 19–23, when performance had stabilized, averaged across subjects. For both $M = 2$ and $M = 4$ conditions, the mixed condition shows the pattern we have come to expect for controlled search, with performance worse at spacing 1 and worse for different targets (NI).

Performance in the categorical condition was equivalent when $M = 2$ and $M = 4$. The usual performance impairment at spacing 1 occurs for identical targets (II), but when the two targets differ (NI), performance in the spacing 0 condition is greatly impaired, a result not seen in previous VM conditions. Furthermore, the II and NI conditions show only a small difference except at spacing 0.

The results from the mixed condition conform to the pattern expected for controlled search in three respects; When $M = 4$ performance is worse than when $M = 2$; there is a performance reduction at spacing 1; and performance is worse for NI than II conditions.

The results from the categorical condition were unexpected in some respects but are explicable in terms of an hypothesis that categories were learned and utilized in controlled search. Let us assume that the presence of a known category allows one to compare the category of the memory set to the category of any given display item in a single operation. Then the $M = 2$ conditions should not differ from the $M = 4$ conditions, and in fact both conditions should show performance levels equivalent to those expected for a memory-set size of 1 in a normal mixed condition. This reasoning explains why the categorical conditions are both superior to the best mixed condition, which uses a memory-set size of 2.

The category hypothesis, however, suggests that target similarity should not matter, so that the II and NI functions should be identical. To explain the spacing 0 results without abandoning the category model, we suggest the following hypothesis. Suppose that when a subject locates a target category, he or she briefly switches to an item mode, perhaps to check that the input is truly a category member. If the second item in the frame is identical, then it will be found in an item-comparison mode, but if nonidentical, it will be located only if the subject reverts quickly enough to a categorical-comparison mode. By the next frame, the reversion to a categorical mode is complete, so the various functions tend to converge. Why would subjects tend to switch to an item mode? The categories were constructed so as to be extremely confusable with each other. Perhaps category encoding is learned under these circumstances but remains somewhat error prone. Then it would be logical to check any target category by using an item mode.

Whatever the explanation for the details of the spacing function, the data as a whole make a strong case that categories have been learned in the categorical condition during VM training, and that the presence of categories in a VM situation allows the subject to adopt a simpler and more efficient form of controlled search.

The argument that controlled search is operating in these conditions would be substantiated if it could be demonstrated that performance improves when the subjects switched to CM training. We turn next to the CM procedure.

3. *Method*

Following the 24th session of VM training, a CM procedure was initiated. In the categorical condition, one of the categories was fixed thenceforth as the memory ensemble, and the other category was fixed thereafter as the distractor set. In the mixed condition, a set of 4 was chosen and was fixed thereafter as the memory ensemble—the remaining 4 items were fixed thereafter as the distractor set. The two sets of 4 used in the mixed condition were those indicated by the position of the comma in the listing of the two sets of eight consonants: {GMFP, CNHD} and {RVJZ, BWTX}. In other words, these were the same sets that had been used as categories in the categorical conditions for other subjects. This CM training procedure was identical for both conditions and followed in other respects the procedures used in the preceding VM conditions. Each subject was given at least 10 sessions of VM training. The frame time after CM Session 4 (Session 24 overall) was reduced to 160 msec and after CM Session 11 (Session 35 overall) was reduced again to 120 msec (at the time of this writing this study had not been completed).

4. *Results and Discussion of the CM Conditions*

The changes over sessions following the switch to CM training are depicted in Figure 5. Consider first the results of Sessions 25 through 28, which immediately followed the switch to CM training. Quite clearly, a very rapid and dramatic rise in performance took place, and furthermore, the results of the mixed and categorical conditions tend to converge. By Session 28 the performance level in the mixed condition had risen to over 90% and in the categorical condition had risen to over 95%. Since the subjects were clearly approaching ceiling, the frame time was then reduced to 160 msec and CM training continued for six additional sessions. Note that performance was still improving and that the mixed and categorical, and $M = 2$ and $M = 4$, conditions effectively converged.

It is interesting to note that during VM training, 20–25 sessions were necessary before the $M = 2$ and $M = 4$ performance levels became equal in the categorical condition. Thus one might tentatively conclude that category encoding in the present experimental context required 20–25 sessions to develop. Thus category encoding would be unlikely to develop for the mixed condition in just four CM sessions. Yet after only four sessions of CM training, the performance level in the mixed condition rose dramatically and approached that of the categorical condition. This result suggests that automatic detection developed in the mixed condition in the absence of a well-learned category (at least during the first four sessions of CM training). In other words, the learning of a category is apparently not a necessary prerequisite for the development of automatic detection.

Figure 7 shows the spacing functions for the six sessions (30–35) of CM training run at a 160-msec frame time. The results are averaged over subjects and sessions. In addition the $M = 2$ and $M = 4$ conditions are lumped together since they did not differ. Somewhat to our surprise the pattern is almost identical to that seen for the categorical condition during the latter stages of VM training, except that the performance level is much higher. The comparable pattern from the CM conditions of Part I/Experiment 3 was quite different: Detection was about equal in all conditions except for a depressed detection rate when spacing was 0 and when the targets were identical. At spacing 0 the present data clearly show performance to be higher when the targets were identical.

Disregarding the shape of the spacing functions, the large improvement in performance when CM training commenced does suggest that the subjects were learning automatic detection. An hypothesis that reconciles these facts suggests that the subjects were using automatic detection to locate the first target, and were then reverting to controlled search to check the accuracy of the target detection. Some time may be lost before the subjects revert to automatic detection. Thus the advantage of II at spacing 0 would be due to the temporary use of a controlled search. This hypothesis is similar to that used to explain the categorical results in the VM conditions (Figure 6). In both cases, a tendency to recheck the first detected item may have led to an alteration in search strategy. The reason may be the same in both cases: Rechecking of located targets may have been induced by the stimulus sets, which were chosen so that the memory set and distractor set were maximally visually confusable. Thus, both automatic detection and category encoding may have been somewhat error prone, thereby requiring rechecking in an item mode. We are currently exploring this hypothesis.

In summary, a number of conclusions can

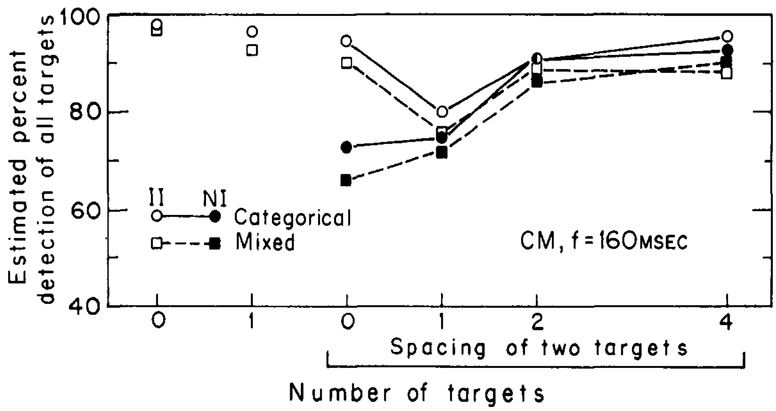

Figure 7. Data from Experiment 3. Estimated percentage of detection as a function of number, spacing, and similarity of targets for CM Sessions 30–35, averaged over memory-set size (which had no effect) for all conditions. (CM = consistent mapping; f = frame time; II = identical targets; NI = nonidentical targets.)

be drawn from the results of Experiment 3. First, arbitrary collections of characters can be learned as categories. For this learning to occur, it is necessary that the categories remain consistently defined across trials, but not necessary that the same category remain the target set across trials. Second, the learning of a category in a VM search paradigm enables the subject to adopt a more efficient controlled search, one in which the category of the memory set may be compared in a single operation to the category of a display item, thus eliminating the effects of variations in memory-set size. Third, a switch to a CM training procedure results in fairly rapid acquisition of automatic detection, with a concomitant improvement in performance for both the categorical and noncategorical stimuli. These conclusions support those suggested by the results of Experiment 2 in all important respects.

D. *What is a Category? A Discussion and Selective Review*

1. *Some Hypotheses Concerning Categories*

Experiments 2 and 3 showed how search benefits from the learning of a categorical distinction between memory ensemble and distractor set. But what is a category? Most generally, we may define any object in memory that refers to (stands for, consists of) any two or more objects in memory as a category. Verbal labels are one common class of categories, but other types of categories might also exist. For example, a visual symbol might represent several objects. Of course, a category in general need not exist in a form similar to any of the sensory modalities. A completely abstract node could represent two or more objects in memory.

With respect to search tasks, the importance of categories is clear. When all of the members of a memory set are members of the same category and no distractors are in that category, then a controlled search can bypass the individual memory-set items and utilize comparisons that involve matching the single category against the category of each display item. For category search to be utilized effectively, however, the category must be learned well enough that each displayed element will be encoded automatically and consistently according to its category. Of course other features, such as the element's name and shape, will also be encoded, but only the category feature is necessary for a category search. We suggest that the category coding must be automatic, because if it were not, a search of long-term memory would have to be carried out to identify the category; the time taken for such a long-term search would almost certainly wipe out any gains that are due to the reduction in number of comparisons allowed by the category search (at least for small memory-set sizes).

Note that automatic category encoding is not the same as automatic detection. Automatic detection refers to the case when a stimulus gives rise to an automatic-attention response that bypasses the need for a serial search through either memory or the display. Automatic category encoding refers to the case when the stimulus gives rise automatically to a node representing the category of the stimulus. However, without an attention response attached to the stimulus node or the category node, the presence of the category would have to be deduced through a serial search of the displayed elements.

Our studies have shown that an automatic category encoding of arbitrary collections of characters (consonants, in the case of Experiment 3) can develop. The cause of the learning is not yet clear. Some subjects in early versions of Experiment 3 never showed any evidence of categorical search (usually subjects with low levels of performance). It may be that these subjects did not notice the categorical nature of the memory ensemble and hence could not develop a well-learned category node in memory. Alternatively, these subjects could have developed an appropriate category node but have failed to use this node to facilitate their search.

There are several lines of evidence suggesting that automatic detection may develop faster if a categorization is available at the start of consistent training. For example, it is easy to learn to search for a set of stimuli defined by a simple physical feature (see Neisser, 1963, 1967, and studies in the next section). The CM results in Figure 5 show that the categorical condition retains some superiority over the mixed condition for about eight sessions (25–32) but the magnitude of the difference is surprisingly small. It is possible that the high confusability between the two categories of Experiment 3 makes the categories and the automatic response difficult to learn. If so, more training prior to the CM phase could possibly have led to a larger difference in the rate of development for the two groups. Alternatively (or additionally) the small size of these categories may have allowed automatic responses to the individual members to develop so quickly that advantages due to a category response were minimized.

It is conceptually important to keep in mind the possibility that automatic responses might be attached to different stages of encoding of a single stimulus. In particular, when an automatic-attention response develops for a category, other attention responses may develop at a different rate for the stimuli making up that category. Thus, if a categorization is available at the start of CM training, the attention responses to the category might develop sooner than attention responses to some or all of the individual stimuli in the category. This effect would be expected since any one stimulus would appear only occasionally as a target, whereas every target would be an instance of the category. On the other hand, in the cases in which the memory ensemble does not form a category at the start of CM training, there might be at least some individual stimuli that come to elicit automatic-attention responses while the category is still being learned.

Once an automatic-attention or encoding response develops, we assume that it is no longer under control of the subject and will occur whenever its corresponding stimulus is presented. This assumption will be supported by Experiment 4. However, when a categorial encoding (but not an attention response) is learned, it will not necessarily be utilized by the subject during a controlled search. We suppose a category response will facilitate controlled search only if the subject both notices the category and also decides to alter his search to compare the category rather than the individual stimuli. These considerations would no longer apply if an automatic-attention response were learned in response to the category; in such a case subject strategy would not matter since attention would be directed to the category in any event.

2. *Search Studies Using Categories*

In a number of memory-search studies (in these studies $F = 1$ and M varies), the memory set consists of several categories, and the distractors might be drawn from these same categories or from different (unknown) categories. The results of these studies all demonstrate the use of some sort of controlled search, perhaps in conjunction with simultaneous automatic detection, that is facilitated by the presence of categories.

For example, Naus, Glucksberg, and Ornstein (1972) and Naus (1974) used memory sets consisting of words from several categories and a distractor set made up of words from these same categories. The results showed linear memory-set-size functions, parallel for negative and positive display items. However, there was a slope reduction when the number of categories increased. The authors suggested that categories were successively searched in random order, each in serial, exhaustive fashion, until the category of the displayed item was reached. When the search of this category was completed, a response was initiated. Homa (1973) presented a related paradigm; his results suggested that a serial, exhaustive search of categories was followed by a serial, exhaustive search within the category of the displayed item. Williams (1971), who presented another related paradigm, also argued that a serial, exhaustive search of categories was followed by a serial search within the category. Okada and Burrows (1973) used multiple categories in the memory set and sometimes cued the subject as to the category of the displayed item in advance of the trial. Their results showed that the search could be restricted to the cued category. Clifton and Gutschera (1971) used memory sets consisting of two-digit numbers and showed that on some trials the subjects would first compare the 10s digits and then would compare the 1s digit only if a 10s digit match was found. Lively and Sanford (1972) used a memory set from one category and a distractor set consisting of some members of that category and other members outside that category. Negative items outside the category of the memory set showed a set-size function with reduced slope.

The various findings in these memory-search studies differ in many details that will not be discussed here. They all tend to show the use of some sort of controlled search that is facilitated by the presence of categories in the memory set.[3]

The final studies to be considered are those visual search studies in which $M = 1$ and F varies, but in which the visually presented items fall in two or more categories, one of which matches the category of the memory-

[3] We shall not attempt at this time to explain why a particular controlled search strategy is adopted in a particular paradigm. Procedural differences among the studies would make such explanations pure guesswork. Another class of studies using categories also deserves mention for completeness, but is not discussed in the main text because of a difficulty in ascertaining whether the effects found in these studies were due to automatic or controlled processing. These are visual search studies ($M = 1$, F varied) that have used memory ensembles and distractor sets from different categories. They have shown either flat set-size functions or curvilinear set-size functions with reduced slopes (e.g., Brand, 1971; Ingling, 1972; Jonides & Gleitman, 1972, 1976). These studies have utilized CM designs with low practice levels; this makes it difficult to ascertain whether the slope reductions are due to the presence of automatic detection, to beneficial effects of categorization on controlled search, or to some combination of these factors.

Note that the slope reductions for visual set-size functions normally should be taken as evidence of automatic detection. That is, in a controlled search one would expect a serial comparison process to be needed to compare the memorized category against the category of each displayed item. However, it was seen in our VM data from Part 1/ Experiment 2 that a reduction of slope occurred when M was equal to 1 and F varied, and we suggested that a "controlled parallel search" might be responsible for this result. A similar argument might be made to explain the slope reductions in the above visual search studies. We do not particularly favor such an explanation because it does not explain why normal set-size functions are found in visual search studies that do not use categories (e.g., Atkinson, Holmbren, & Juola, 1969).

A finding related to those from studies using categories in visual search is that of De Rosa and Morin (1970) in a memory search study. They showed that when the memory set consisted of consecutive digits, then both the positive and negative reaction times increased as a function of the numerical closeness of the displayed item to the boundaries of the consecutive memory set. The CM procedure was used with small amounts of practice, so it is again difficult to ascertain whether such effects were due to the action of automatic detection or to some controlled search facilitation allowed by the categorical character of the memory set.

set item. Smith (1962) and Green and Anderson (1956) carried out similar studies in which a field of colored two-digit numbers was presented and the subject was asked to search for a particular two-digit number in a particular color. The results showed that search rate depended much more on the number of items in the designated color than on the total number of items presented, though both effects were present. A VM procedure was utilized in their studies so that automatic detection could not have been used to restrict search to the items of a particular color. Probably a two-phase controlled search was utilized in this situation. First, a rather fast serial search was probably carried out to locate the next item of the designated color; then a slow comparison of that two-digit number was probably made to determine whether it was the target. This explanation is supported by the fact that searching took a long time (up to 20 sec), so that a 30–40 msec search rate for color could have accounted for the increases in response time as total size increased. These increases were small only in comparison with the larger within-category effects.

The main conclusion to be drawn from the studies reported in this section, based on repeated findings, is that the presence of categories in search tasks can be used to facilitate, benefit, and modify controlled search.

III. Experiments on Focused Attention

In part I we made the case that the processes utilized in attention experiments and those utilized in search and detection experiments are often the same. In fact, in many cases it is purely arbitrary whether a given study is referred to as an "attention," "search," or "detection" study. Experiment 1 in Part I could be described as a divided-attention study in which attention had to be divided among M memory-set items and F frame items during each frame. The results showed tremendous deficits in dividing attention in the VM conditions, and virtually no deficit in dividing attention in the CM conditions. The results of Experiment 2 of Part I, and the fit of the quantitative model to both studies, showed that divided-attention deficits in these paradigms are due to the limitations of controlled processes, in particular, to the limited rate of short-term search.

Thus, all our comments and conclusions concerning search and detection apply equally well to attention studies. In particular, CM training should lead to the development of automatic detection that shoulld bypass divided-attention limitations, while VM training should cause controlled search that should severely limit the ability to divide attention.

In discussing the development of automatic detection we have proposed that an automatic-attention response is learned in response to the unchanging members of the memory ensemble. The present study will test this hypothesis. The test (Experiment 4d) will entail asking the subject to ignore certain locations and then inserting in those locations items that subjects had been previously trained to respond to as CM targets (to see whether these targets attract attention).

These studies will also answer the following question: To what degree can the subject focus attention on a specified subset of the inputs without distraction from the remaining (irrelevant) inputs. Such studies are usually termed *focused-attention* studies, in contrast to the studies of Part I and Experiments 1 to 3 of the present article, studies that would be appropriately termed *divided-attention* studies.

A. *Terminology*

There is a problem of terminology in the studies of this section that is best solved by the introduction of the following definitions:

1. A *foil* refers to any input that appears in a to-be-ignored display location, whereas a *distractor* refers to a nontarget that appears in to-be-attended display location.
2. A *CM foil* is a foil that has previously been used in CM training as a memory-set item.
3. A *CM target foil* is a CM foil that would have been a target requiring a positive response if it had appeared in a to-be-attended

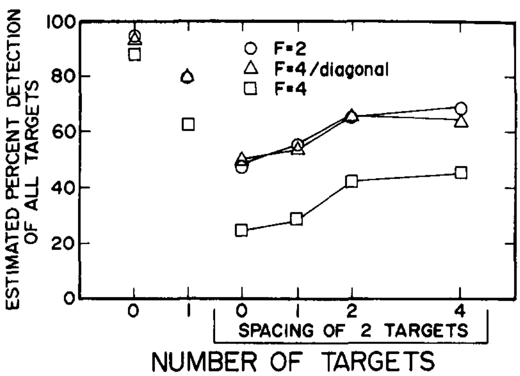

Figure 8. Data from Experiment 4a. Estimated percentage of detection as a function of number and spacing of targets in a VM procedure with frame time equal to 200 msec. The spacing 0 points are nonidentical targets, while the spacings of 1, 2, and 4 are identical targets. (F = frame size.)

display location (although it appears in fact in a to-be-ignored display location). It is a member of the current memory set.

4. A *VM foil* is a foil that has previously served as both a target and distractor in VM conditions.

5. A *VM target foil* is a VM foil that would have been a target requiring a positive response if it had appeared in a to-be-attended display location. It is a member of the current memory set.

6. *Valid* positions or characters are to-be-attended positions or characters. *Invalid* positions or characters are to-be-ignored positions or characters.

7. *FII* denotes a trial in which two identical memory-set items appear, one in a to-be-attended display position, and one in a to-be-ignored display position.

8. *FNI* denotes a trial in which two different memory-set items appear, one in a to-be-attended display position, and one in a to-be-ignored position.

B. *Focusing Attention in a VM Multiple-Frame Task*

The first study in the present series of experiments is designed to show that "controlled search" is not a misnomer, that subjects can control their search to the extent that VM foils can be ignored, and that search can be carried out through the valid display positions without decrement caused by the foils.

1. *Method*

The paradigm utilizes the multiple-target, multiple-frame VM procedure. There are three main conditions: (a) $M=2$, $F=2$; (b) $M=2$, $F=4$; (c) $M=2$, $F=4$, diagonal. Condition (c), denoted "diagonal," was designed so that one of the diagonals of the display was always valid (to-be-attended), and the other diagonal was always invalid (to-be-ignored). In this condition four characters were presented on each frame, two valid, and two invalid foils. We expected that this diagonal condition would elicit performance similar to that for the $M=2$, $F=2$ condition, and better than that of the $M=2$, $F=4$ condition.

The four subjects in this study were the same as those used as in Experiment 3 of Part I. In the diagonal condition, only the upper-left and lower-right frame positions ever contained a target. The subjects were fully instructed regarding this fact and were instructed to ignore the invalid diagonal. For each subject, the blocks of trials utilizing the diagonal procedure were run after the other conditions were completed. The $M=2$, $F=2$ condition had 14 blocks per subject; the $M=2$, $F=4$ condition had 11 blocks per subject; the diagonal condition had 16 blocks per subject, the first 4 of which were practice blocks. There were 120 test trials per block, plus an additional 30 trials of practice for the first block of a session and 15 trials of practice for each subsequent block in a session.

The procedure used in this study differed somewhat from that of Part I/Experiment 3 and from the other multiple-target tasks reported in this article. The primary difference lay in the relation of target similarity to spacing. The spacing 0 condition utilized NI targets only, and the spacing 1, 2, and 4 conditions utilized II targets only. In addition, the present experiment allowed targets to reoccur in the same frame positions if they were not in successive frames. Also, in the present study, the subjects pushed a single response button each time they thought they detected the target. The responses were to be made when the targets appeared rather than at the trial's end.

2. *Results and Discussion*

On about 1% of the trials in each condition a response was made before the target appeared, and these responses were not counted. The results are presented in Figure 8. The results of the $F=2$ and diagonal conditions obviously do not differ, but results of both conditions are clearly superior to those of the $F=4$ condition. These results were expected on the supposition that the subject should have been able to control his processing in the diagonal condition so that comparisons would occur only for the characters on the valid diagonal.

The peculiar relationship of spacing to target similarity in this study caused the shape of the spacing functions to appear to differ from those in previous multiple-target studies. In fact, however, the corresponding points from the present study and Part I/ Experiment 3 are quite close in value. It is interesting to note the subject's comments at the end of the experiment. They noticed an excess number of II conditions overall, but none noticed that the II pairs were not tested at spacing 0, nor that the NI pairs were not tested at the longer spacings.

The implications of this study are straightforward. Subjects can control their search in VM situations at least to the degree that comparisons can be limited to a specified diagonal. It is possible that characters on the invalid diagonal are sometimes compared, but not until the valid diagonal is searched first (if time were taken to compare foils during the search of the valid diagonal, then performance would be worse). Of course, this study demonstrates only a minimum amount of subject control. In future studies it would be desirable to explore this matter more thoroughly. For example, we might ask: Can search alternate between diagonals in successive frames? Can search order be cued individually for each frame?

C. *Distraction Caused by "Targets" During VM Search*

Experiment 4a showed no distracting effect of VM foils (on the invalid diagonal). However, none of these foils was in fact identical to any member of the memory set. The present study, then, is designed to determine the distracting effect of VM target foils: To what extent can members of the memory set be ignored when they appear in invalid display locations?

1. *Method*

The paradigm is similar in general outline to that of the VM conditions of Part I/Experiment 3. The peculiarities of Experiment 4a were eliminated. The multiple-frame procedure was utilized with $M = 2$ and $F = 4$. The upper-left, lower-right diagonal was valid, and the other diagonal was invalid. Either zero or one targets appeared on the valid diagonal, and an appropriate binary response was required. A VM task was utilized on the valid diagonal, and the foils on the invalid diagonal were chosen from the distractor set for the valid diagonal. In addition, there was on every trial exactly one VM target foil (i.e., a member of the memory set) on the invalid diagonal.

On one-third of the trials, only a VM target foil appeared, always on frames 8–13. On two-thirds of the trials both a target foil and target appeared; on one-half of these trials the target and target foil were identical (FII), and on one-half of these trials the target and target foil were different (FNI). Finally, the target appeared on any of frames 8–13, and the target foil appeared equally often on frames −4, −1, +1, and +4 with respect to the target frame. That is, the spacing between the target and target foil was systematically varied. Thus, the target-foil-only trials occurred with a probability of $\frac{1}{3}$, and each of the other eight conditions occurred with a probability of $\frac{2}{3} \times \frac{1}{8} = \frac{1}{12}$. All these trials were randomly intermixed within each block.

Each block contained 144 test trials preceded by 15 practice trials. Each subject ran through a total of 12 or 13 blocks in four sessions.

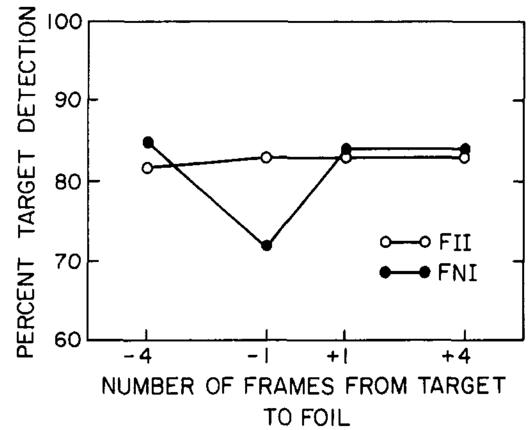

Figure 9. Data from Experiment 4b. FII = target foil identical to target; FNI = target foil nonidentical to target. Percentage of target detection as a function of the spacing and the similarity between the target and the target foil. Varied-mapping procedures were used, and frame time was 200 msec.

2. *Results and Discussion*

The results are shown in Figure 9. The figure gives the percentage of hits as a function of the spacing between the target and the target foil, when the target foil is identical (FII) and nonidentical (FNI) to the target. The correct rejection rate when no target was present was 91%. The results may be summarized briefly: The target foil had no effect on the hit rate except when it was different from the target (FNI) and preceded the target by 1 frame, in which case detection probability was decreased by .11.

In the present study a target foil reduced performance when it preceded, but not when it followed, the target, Suppose that this reduction is due to the same factors that produce decrements in multiple-target detection in VM conditions (Experiment 3, Part I). Then one target probably inhibits detection of a subsequent target, but not of an antecedent target. That is, the reduced detection at spacing 1 seen in the VM conditions of Part I/Experiment 3 and also seen in Figure 4 and 6 of the present paper, could have been caused by either of the two targets' reducing detection of the other. The present results provide some reason to argue that it is the first of two successive targets that reduces detection of the second.

Let us assume that the decrement in percentage of detection of two targets in VM paradigms is in fact a decrement in detection of the second target. Then it is possible to compare the magnitude of the decrements in the multiple-target studies (e.g., Experiment 3, Part I) and the present study. In fact, the decrements are much larger in the earlier studies in which all display locations were valid. For example, the difference in detection between spacing 1 and spacing 4 in Part I/Experiment 3 was 30% in the NI condition and 8% in the II condition; the comparable decrements in the present study are 11% and 0%. Thus a target foil reduces detection of a target in the next frame, but the reduction is much smaller than that caused by an actual target.

How can the present findings be reconciled with those of Experiment 4a in which foils did not impair performance? The pattern of findings would be explicable if the valid diagonal were always searched first and then, whenever that search finished early, one or more characters on the invalid diagonal were inadvertently checked in addition. In this event, target foils would occasionally be noticed and might harm subsequent detection in a fashion similar to that caused by true target detection. Of course, this hypothesis is speculative and other explanations are undoubtedly available.

Whatever the explanations for the details of the results, it is safe to summarize the results of Experiments 4a and 4b as follows: Subjects are able to control their search in varied-mapping paradigms. The degree of their control is sufficient to eliminate any distracting effect of nontarget foils and to reduce the distracting effect of target foils well below the level caused by other targets. Thus, attention focusing is quite successful: VM foils have, at most, a small distracting effect when they appear in to-be-ignored locations during controlled search.

D. *The Distracting Effect of "Targets" During CM Search*

Experiments 4a and 4b examined the ability to focus attention during controlled search. We next ask: To what degree does attention focusing affect automatic detection? Our previous studies have demonstrated that automatic detection is not affected by frame size, so there would be no point in carrying out a CM counterpart of Experiment 4a in which the invalid diagonal would contain distractors only. We decided to carry out a counterpart to Experiment 4b to find out whether a CM target foil would interfere with automatic target detection on the valid diagonal.

M and *F* were set equal to 2. During each frame, each diagonal contained one mask and one character. One diagonal was always valid, the other invalid. On every trial a CM target foil appeared on the invalid diagonal in exactly one frame. A CM target appeared on the valid diagonal on two-thirds of the trials; on one-half of these trials the target and target foil were identical (FII) and on one-half, nonidentical (FNI). In each of these cases, the target foil occurred either four frames before, in the same frame with, or four frames after the target ($\frac{1}{3}$ probability each). There were 162 trials per block and 6 blocks per each 1-hour session. Two sessions were run with frame time equal to 60 msec and two sessions were then run with frame time equal to 30 msec.

The results are shown in Table 2. The only significant effect was a slight drop in performance (about 4%) in the FII condition at $f = 60$ msec with simultaneous target and foil ($z = 4.45$, $p < .0001$); performance dropped considerably when f dropped to 30 msec, but all conditions became roughly equal.

The results of this study demonstrate that a target foil hinders detection of an identical simultaneous target. The magnitude of the effect is small, but so is the magnitude of the decrement that occurs when two simultaneous CM targets are presented. In fact these decrements are not much different at the equivalent frame times (4% vs. 8%). It seems reasonable to assume that the same mechanism is causing both effects.[4]

It is interesting to compare these findings to those of Eriksen and Eriksen (1974). In our terminology, they used a single-frame CM task with $M = 2$, $F = 1$, and they used reaction time as the dependent measure. However, the distractor set and the memory ensemble were both of size 2 and were consistently mapped across trials, so that an equally strong but opposite tendency to automatically respond to the members of the two sets was probably learned. (In our CM studies, only the memory-set items can come to elicit automatic responses, because distractors are presented on every trial. The only exception in our studies occurred when $F = 1$ in the single-frame paradigm of Part I, and even then the distractor set was always larger than the memory set.) In the Eriksen and Eriksen study, the relevant item always appeared directly above the fixation point, but three invalid items were placed on each side of the relevant item. The nature and distance of these invalid items were varied. When the distance to the nearest item reached 1° of visual angle, the invalid characters had little effect on reaction time. At closer distances, all reaction times were slowed, but flanking distractors produced the greatest slowing (since they automatically produced a competing response); flanking characters that were not distractors or memory-set items produced a moderate slowing (regard-

Table 2
Effect of Distraction on Automatic Detection: Experiment 4c

Variable	Spacing (target to foil)	% hits	% correct rejections
Session 1			
60-msec frame time			90.0
FII	−4	91.9	
FNI	−4	94.2	
FII	0	89.4	
FNI	0	95.6	
FII	+4	96.1	
FNI	+4	95.1	
Session 2			
60-msec frame time			91.2
FII	−4	95.4	
FNI	−4	93.5	
FII	0	90.7	
FNI	0	93.8	
FII	+4	92.8	
FNI	+4	95.6	
Session 3			
30-msec frame time			75.2
FII	−4	75.7	
FNI	−4	79.1	
FII	0	74.5	
FNI	0	77.6	
FII	+4	77.1	
FNI	+4	80.1	
Session 4			
30-msec frame time			71.5
FII	−4	77.3	
FNI	−4	78.9	
FII	0	80.3	
FNI	0	81.0	
FII	+4	78.7	
FNI	+4	78.2	

Note. FII = trial in which two identical memory-set items appear, one in a to-be-attended display position, and one in a to-be-ignored display position. FNI = trial in which two different memory-set items appear, one in a to-be-attended position, and one in a to-be-ignored position.

[4] The results in Table 2 give some indications that the disrupting effect may change with practice. Over the four sessions, the differences between FII and FNI at spacing 0 was, respectively, 6.2%, 3.1%, 3.1%, and .7%. If this trend actually exists, it may be caused by the development of a new automatic response that restricts search to the valid diagonal. That is, position-specific information might govern the automatic-attention response; such an automatic process could be learned because the valid diagonal never changes over trials. It would be interesting in future investigations to compare performance in this condition with performance in a condition that alternates the valid diagonal from trial to trial, since alternation should prevent the long-term learning of position-specific encoding.

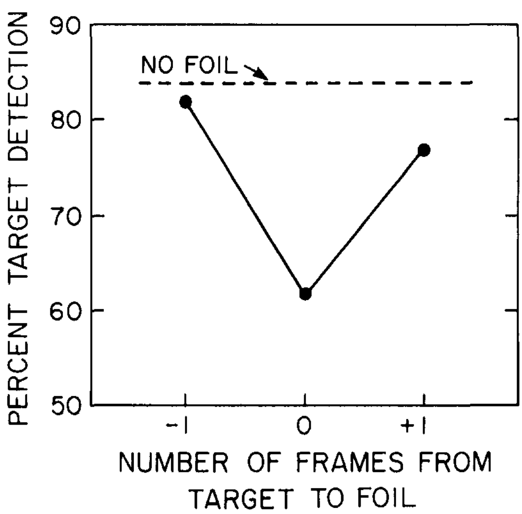

Figure 10. Data from Experiment 4d. Percentage of target detection in a varied-mapping procedure as a function of the spacing between the target and the target foil, when subjects had previously been trained to respond to the foil as a consistent-mapping target. Frame time was 200 msec.

less of their visual similarity to either set), and flanking targets, whether identical or nonidentical, produced the least slowing (since they automatically produced a compatible response).

Both the Eriksen and Eriksen (1974) study and our study suggest that CM target foils can neither be excluded from processing nor prevented from affecting performance. In addition, however, the Eriksen and Eriksen result suggests that the spatial configuration of the inputs can determine the magnitude of these effects (see Footnote 4 for a discussion of location-specific automatic detection).

E. *The Distraction of Controlled Search by Automatic Detection*

Controlled search depends on an apparently serial process that should easily be disturbed if an automatic-attention response occurs that draws attention to an invalid location. Such a rationale suggests a paradigm for our next study in which CM foils appear on the invalid diagonal while a controlled search occurs on the valid diagonal.

This study is probably the most important in this series (4a–4d) because it provides a test of the hypothesis that an automatic-attention response to consistently trained targets is learned.

1. *Method*

Each frame contained 4 characters. The upper-left to lower-right diagonal was valid. Memory-set size was 2, and the VM conditions were utilized. Search on the valid diagonal was for consonants in consonants (or digits in digits for other subjects). The foils on the invalid diagonal were chosen from the distractor set used for the valid diagonal, except for the CM foil. When a CM foil appeared it was chosen from the set that had served as the CM memory ensemble in all previous studies for that subject. Thus, if consonants were being used on the valid diagonal, a CM foil would be a digit, and vice versa.

On one-sixth of the trials, neither a target nor a CM foil appeared; on one-sixth of the trials only a target appeared; on one-sixth of the trials only a CM foil and no target appeared; on one-half of the trials both a target and a CM foil appeared, with the spacing between them equally likely to be −1, 0, +1. Each block of trials contained 144 test trials. The first block of each session had 15 practice trials and the other three blocks had 5 practice trials. Two sessions were run for each of the four subjects (the same subjects as those used in Experiments 4a–4c.) The frame time was 200 msec.

2. *Results and Discussion*

The results are shown in Figure 10. Each percentage is based on 768 observations. When no CM foil or target appeared the false alarm rate was 4%. When only a CM foil appeared the false alarm rate was again 4%. Thus, in the absence of an actual target the presence of a CM foil caused no increase in false alarms. Since the appearance of CM foils was highly correlated with the appearance of targets, one might have supposed that a bias or guessing strategy would arise, such that subjects would more often guess that targets were present if a CM foil was seen. The present data argue against such a bias (as do the data discussed below).

The hit rate when a target but no CM foil was present was 84%. The distraction caused by a CM foil may be determined through comparison with this performance level. When the CM foil *preceded* the target by one frame the hit rate was 82%, not significantly lower than the 84% base rate. When the target and CM foil were presented simultaneously, the hit rate dropped to 62%. And when the CM foil *followed* the target by one frame the hit rate rose to 77%, a figure still significantly lower statistically than the 84% hit rate when no CM foil was present. In short, a CM foil provides little hindrance to the detection of a VM target presented in the following frame, greatly hinders detection of a VM target presented simultaneously, and even reduces slightly the detection of a VM target presented in the preceding frame.

The simplest interpretation of these results is that a CM foil interrupts the ongoing controlled search and causes a loss of processing time, thereby reducing target detection. The interruption is caused by what we have termed an automatic-attention response to the CM foil. We have argued that subjects have been trained to respond to CM targets with an automatic-attention response. Even when such a target appears on a to-be-ignored display diagonal, it apparently causes an attention response that interrupts processing along the valid diagonal and directs attention to the invalid diagonal. The time lost before attention is returned to the valid diagonal and the search is resumed causes a considerable decrement in performance if the target is in fact on the valid diagonal during that frame.

Since target detection in a frame following a CM foil is not hindered, the recovery from the distraction caused by a CM foil must be quite rapid; that is, recovery must take place in a time period under 200 msec. Furthermore, the automatic-attention response must occur fairly rapidly, since a CM foil even reduces detection of a target in the preceding frame. Perhaps the processing of one frame occasionally overlaps the start of the next; for example, a comparison initiated but not completed before termination of a frame may be completed before controlled processing of the next frame begins. Occasionally, a target might be the character undergoing comparison when the subsequent frame begins, and a CM foil in that next frame could interrupt comparison of the target and thereby impair performance.

The primary finding of the present study is the demonstration that the responses to CM targets are both well learned and automatic. CM targets cannot be ignored, even when they are known to be irrelevant and occur in consistently invalid display locations and even when subjects are instructed to ignore them.

Our results, of course, do not establish that it is impossible to control the detection processes that we have labeled "automatic." In fact, Sperling (Note 1, Note 2) has reported some findings that do tend to show that at least some amount of control over some aspects of the search process in CM multiple-frame tasks is possible. However, the degree of control over, and the ability to ignore, stimuli that subjects have been trained to see as targets in CM situations are clearly much less than in VM situations. Thus, for the processing mode in the CM conditions of the present tasks, we prefer the possibly slightly inaccurate, but descriptive, term *automatic* to the logically more accurate, but less descriptive, term *systemic*, which was used in Shiffrin (1975a) to refer to the same type of processing.

F. *Limitations on the Ability to Focus Attention*

The present series of studies (4a–4d) examined the subjects' ability to focus attention. Experiment 4a showed that attention during controlled search could be focused on specified spatial locations. Experiment 4b showed that the focusing is not complete, because certain types of stimuli in to-be-ignored locations are sometimes processed and when processed, can reduce target detection. Experiment 4c showed that a CM target in a to-be-ignored location hinders automatic detction by a small amount. Experiment 4d showed that attention is greatly distracted from the relevant locations when a CM target appears in a to-be-ignored location during a VM search task. This last finding is particularly interesting because the distracting stimulus differs from the relevant stimuli not only in spatial location but also in category.

These findings are particularly important because previous demonstrations of difficulties in focusing attention have usually utilized incompatible responses. That is, a stimulus in response to which a subject has been trained to emit response x is presented in a to-be-ignored location in conjunction with a stimulus requiring a response y that is incompatible with x. In the Stroop Color–Word Interference Test (1935), for example, the subject must read names of colors that are printed in ink colors that do not correspond to the names. Reading speed is greatly slowed in this case (see Jensen & Rohwer, 1966). Keele (1972) showed that decrements in processing stimuli resulted when a similar color–name incompatibility was present. The Eriksen and Eriksen (1974) study discussed above also showed that response time was dependent on the compatibility of responses between stimuli to be attended to and those to be ignored.

Many other examples of this kind can be found in the literature, but our demonstrations differ in one very important respect. Namely, the distracting stimulus in our studies hurts performance even though subjects have been trained to respond to it in exactly the same way as to the target stimulus. Both types of studies demonstrate that the to-be-ignored stimulus receives processing, but the studies utilizing incompatible responses do not demonstrate any processing interaction short of the motor responses themselves. On the other hand, our studies suggest that the focused-attention deficit results from a diversion of attention and an accompanying loss of processing time, since there is no response incompatibility.

Our studies have shown that focused-attention deficits can arise due to distraction by either CM or VM items in appropriate contexts, though probably due to rather different mechanisms in the two cases. In many studies in the literature that have shown focused-attention deficits, it is not clear which factor is responsible. For example, in shadowing tasks the listener repeats aloud a designated message (usually in one ear) while other distracting messages are also presented. The early studies showed that distracting messages barely harm shadowing performance when they are distinguishable by an obvious physical characteristic (such as spatial origin; see Cherry, 1953; Cherry & Taylor, 1954). These studies are probably like our Experiment 4a in which one diagonal is relevant during VM search and the other can be ignored. A result related to Experiment 4b can be found in studies by Treisman (1964a, 1964c) in which items of some possible relevance are presented in the nonshadowed ear (similar to our target foils). For example, if the distracting message is in the same voice, but in French, a considerable distraction occurs. Many studies have shown that the message in the nonshadowed ear is analyzed to some considerable depth (e.g., Lewis, 1970; Treisman, 1960, 1964b). In these studies it is not clear which type of processing causes the analysis, but in other studies it is clear that automatic processing is responsible for the analysis that occurs (Corteen & Wood, 1972; Von Wright, Anderson, & Stenman, 1975).

Further reviews of such studies will not be undertaken here because of the difficulty in ascertaining whether focused-attention deficits are caused by controlled processing of invalid items or by automatic processing of invalid items.[5] If nothing else, our results strongly suggest that future research on focused attention should include controls to differentiate deficits caused by controlled and automatic processing.

In summary, we have seen that subjects can divide attention almost without deficit when automatic detection is utilized (Part I/

[5] Controlled processing might be given to occasional inputs in an irrelevant ear, because ear of origin is a fairly difficult cue to utilize (compared with spatial location in a visual task, for example). Whether automatic processing or controlled processing is used to restrict analysis to the desired ear, occasional failures of selection are to be expected, failures that will cause occasional items in the to-be-ignored ear to be given controlled processing. In addition, controlled processing of items in the to-be-ignored ear may occur on occasion when processing of the items in the valid ear is completed and a brief interval occurs before the next valid input arrives (an argument similar to this is used in the discussion of Experiment 4b).

Experiments 1 and 2, CM conditions, see Figure 2) and cannot divide attention without deficit when controlled search is utilized (Part I/Experiments 1 and 2, VM conditions, see Figure 2). Focused-attention deficits are quite substantial when caused by automatic responses to to-be-ignored items (Experiment 4d and Figure 10) but are quite a bit smaller when caused by controlled processing of to-be-ignored items, since the subjects' control is usually adequate to prevent such processing (Experiments 4a and 4b, and Figures 8 and 9). One implication is the prediction that the type of training procedures, VM or CM, will determine whether divided- or focused-attention deficits will occur.

G. *The Nature of Automatic Detection*

The reversal studies, Experiments 1 and 2, and the category study, Experiment 3, make it clear that long-term learning is responsible for the phenomenon of automatic detection that is seen in the CM conditions. The lengthy training period required for acquisition, and the even longer training periods required to alter the detection process once learned, testify to the permanent nature of the learning. Furthermore, the phenomenon is powerful enough to transcend the laboratory context. Subjects given CM training on one group of letters as memory-set items with another group as distractors, as in Experiments 1 and 3, report effects on reading outside the laboratory. Despite the fact that all the letters used in the experiments were capitalized, the subjects reported that the memory-set items from the CM conditions tended to "jump out" from the page during normal reading. This effect was distracting enough that one subject would not attempt to read for an hour or more after an experimental session.

It is clear then that automatic detection reflects a powerful long-term process. We now ask: What is it that is learned? Experiments 4a–4d imply that an attention response is learned, but many details of the automatic detection process remain to be specified.

It is our feeling that several factors contribute to the process that has been termed automatic detection. First there is an automatic-attention response to the features that are encoded from an input target; this response directs attention to the representation in short-term memory of the relevant visual input and also to the representation in memory of the appropriate member of the memory set. Second, in addition, an automatic "target" response is learned that tells the subject that a target is among the inputs. Third, in addition, in some situations an automatic overt motor response (such as a button press) is learned in response to a target.

The third factor, an automatically produced overt motor response, is probably limited to special situations in which the same completely consistent response to all targets is immediately required. Our reaction time study (Part I, Experiment 2) may represent such a situation, but our multiple-frame tasks do not. In the multiple-frame tasks, no overt response is required until the trial's end. In addition, some of the tasks in our studies require that the number of targets be counted, rather than that a response be made to each target as it appears. Finally, the tasks in Experiment 4 occasionally present target foils in to-be-ignored locations, and the subject has little difficulty suppressing responses to such targets. Thus, we feel that automatic overt responses can be learned but are probably not a major contributor to performance in most of our tasks.

The first two factors, the occurrence of automatic "attention" and "target" responses, undoubtedly play a major role in our tasks. However, such responses must be followed by some sort of controlled process or controlled decision to generate the required overt response. A number of controlled processes may be used. In most of our studies it would not have been sufficient to simply use the *occurrence* of an automatic "attention" or "target" response as a basis for a decision to respond. For one thing, some of our studies require responses to be counted; for another, some of our studies present target foils to which responses must be withheld. Furthermore, in Experiment 3 we saw that the memory and distractor sets could be so confusable that automatic responses could not be learned in a perfectly accurate manner, so that a switch to a controlled item mode was necessary to check the accuracy of automatic responses.

For all of these reasons, we propose that following the occurrence of an automatic-attention or target response, the subject engages in the minimal controlled processing necessary to satisfy the task requirements. Such controlled processes might consist of comparing the memory representations to which attention has been drawn, counting the target instances, and checking the spatial location of located targets to see whether they are foils or valid targets.

Before concluding this discussion, it is useful to consider briefly a few supplementary hypotheses regarding the automatic detection process. Each of these hypotheses assumes the occurrence of an automatic-attention response.

First, consider the possibility that attention is drawn automatically to the representation of the relevant visual input, which is then compared serially to the memory set. This hypothesis is easily rejected since it implies, contrary to fact, that memory-set size will have a large effect under CM conditions.

Second, consider the possibility that an attention response directs the subject to the relevant visual input, whose *category* is then compared in one step to the *category* of the memory set. This hypothesis is testable in any of several ways but is difficult to reject on the basis of present evidence. A tentative inference from Experiment 3 suggested that automatic detection for items in a four-letter set developed much faster than category learning for that set. If so, then this category hypothesis could be ruled out (because the $M = 2$ and $M = 4$ CM functions converged quickly in the mixed condition of Experiment 3). Nevertheless, additional research will be needed to test this hypothesis conclusively.

Assuming that an automatic-attention response underlies automatic detection, it seems clear why automatic detection does not develop in VM situations: An attention or overt response that is helpful on one trial (when the producing stimulus is a target) will be harmful on another (when the producing stimulus is a distractor). A subject cannot learn both an attending and a nonattending response to the same stimulus. While this argument seems clear, an important theoretical question remains. What is the underlying mechanism that inhibits the learning of an automatic-attention response in the VM search situation?

Suppose that a learning event takes place each time a target is found correctly and is therefore reinforced. In consistent (CM) paradigms there will be no impediment to the learning of attention responses. In VM paradigms, the outcome is less clear. It might be supposed that every item, whether currently a memory-set item or a distractor, comes to elicit attention responses due to those trials on which it is a target. There are then several possibilities: (a) Learning continues until all items have acquired attention responses of roughly equal strength. Because the responses are of equal strength, they tend to cancel each other, and controlled search must be utilized. Later, if a switch is made to a CM paradigm, the strength of the responses to the fixed memory-set items becomes greater than that for the distractor items. (b) As an attention response begins to develop it starts to occur on trials when the input is a distractor. Because attention is directed to the wrong input, performance suffers, and the response is therefore inhibited. (c) Each time a distractor is compared during a controlled search, whether or not an attention response occurs, inhibition of any previously reinforced attention response may take place because a comparison is carried out but not reinforced. According to explanations (b) and (c), the inhibition cancels any learning that would otherwise occur, so that attention responses to any of the items never develop to significant degrees.

We prefer hypotheses (b) or (c), or any similar proposal that implies that attention responses do not develop in VM paradigms. Is there any evidence to distinguish these views from hypothesis (a)? Only indirect evidence is available at present, but the models are easy to test. For example, after VM training one could introduce new items as targets, with other new items as distractors, or with the previously trained VM items as distractors. The old VM distractors should severely hinder performance if they elicit attention responses developed during previous training. Other similar tests could also be carried out, but a definitive answer is not yet available at the time of this writing.

IV. A Framework for Information Processing

This section of the paper will be organized as follows. Section A will present an overview of the system and an overview of the role of controlled and automatic processing. Section B will elaborate on the fundamental nature of controlled and automatic processing. Section C will present a framework for search, detection, and attention. Section D will discuss how automatic and controlled processing are utilized in memory storage and retrieval.

A. *An Outline of a General Theory*

Memory is conceived to be a large and permanent collection of nodes, which become complexly and increasingly interassociated and interrelated through learning. Most of these nodes are normally passive and inactive and termed *long-term store,* or *LTS,* when in the inactive state. The set of currently activated nodes is termed *short-term store,* or *STS.* LTS is thus a permanent, passive repository for information. STS is a temporary state; information in STS is said to be lost or forgotten when it reverts from an active to an inactive phase. Control of the information-processing system is carried out through a manipulation of the flow of information into and out of STS. These control processes include decisions of all sorts, rehearsal, coding, and search of short- and long-term store. LTS contains learned sequences of information processing which may be initiated by a control process or by environmental or internal information input, but are then executed automatically with few demands on the capacity of STS.

1. *Long-Term Structure*

The structure of the nodes making up LTS will be treated as a very general graph with complex interrelations among nodes. An individual node may consist of a complex set of information elements, including associative connections, programs for responses or actions, and directions for other types of information processing. What then sets off one node from a group of nodes? One node is distinguishable from a group of nodes because it is unitized, that is, when any of its elements are activated (i.e., placed in STS), all of them are activated. One activated node may of course activate another node, but it does not do so in all situations, only when the context or the state of the information-processing system is appropriate.

2. *Structural Levels*

It seems likely that the structure of long-term store, at least in part, is arranged in levels (perhaps sometimes in a hierarchical tree structure). By *levels* we refer to a temporal directionality of processing such that certain nodes activate other nodes but not vice versa. In sensory processing, there is a tendency for increasing information reduction as successive levels are activated. Thus, a visually presented word could first be processed as a pattern of contrast regions, colors, regular variations, and so forth; then lines, angles, and other similar features could be activated; then letters and letter names and verbal or articulatory codes; then the word's verbal code; and finally, the meaning and semantic correlates of the word. This sequence is meant as an example, and we do not wish to imply that these are the relevant features, that this is the only possible ordering, or that this listing is exhaustive. Such a sequence of feature encoding should occur automatically to a normally skilled reader.

3. *Automatic Processes*

An automatic process can be defined within this system as a sequence of nodes that nearly always becomes active in response to a particular input configuration, where the inputs may be externally or internally generated and include the general situational context, and where the sequence is activated without the necessity of active control or attention by the subject.

An automatic sequence differs from a single node because it is not necessarily unitized. The same nodes may appear in different automatic sequences, depending on the context. For example, a red light might elicit a braking response when the perceiver is in a car, and elicit a walking, halting, or traffic-scanning response when the perceiver is a pedestrian.

Since an automatic process utilizes a relatively permanent set of associative connections in long-term store, most automatic processes will require an appreciable amount of training to develop fully. Furthermore, once learned, an automatic process will be difficult to suppress or to alter.

When an automatic sequence is initiated, its nodes are activated and hence the associative information enters STS. This fact does not, however, mean that the various elements of the automatic process must be available to consciousness or recallable at a later time. The activation in STS could be extremely brief (in msec, say) and unless attention is directed to the process when it occurs or unless the sequence includes an automatic attention-calling response, then the information may be immediately lost from STS, and the subject may be quite unaware that the process took place. Even when an automatic-attention response is part of the sequence, it will not necessarily affect ongoing controlled processing unless the strength of the attention response is sufficiently high.

4. *Thresholds of Activation*

However an automatic process is initiated, whether by internal or external input or by a control process, it is presumed that the probability that the process will run to completion depends on the strength of the initiating stimulation. The simplest and most common examples are seen in studies of psychophysical thresholds. If a letter, for example, is visually presented at a low-enough intensity or a short-enough duration, then it may be encoded not as a letter but as a partial collection of line-like features. At lower durations or intensities even these features will not be activated.

Although we have described automatic processes as largely beyond subject control, some indirect control is possible through manipulation of the activation threshold for automatic processes. In particular, according to what we shall call the *contextual hypothesis*, the activation threshold can be lowered by the inclusion of information in STS (at the time of presentation of the activating input) that is associatively related to the nodes making up the automatic sequence.[6]

Note that a lowering of the threshold does not imply that the quality of processing is improved. One result of threshold lowering is that the automatic process will be triggered by inputs that normally would and should not do so. For example, the feature "horse" might be incorrectly triggered by the input "house" if the threshold for "horse" is lowered sufficiently.

5. *Controlled Processes*

A controlled process utilizes a temporary sequence of nodes activated under control of, and through attention by, the subject; the sequence is temporary in the sense that each activation of the sequence of nodes requires anew the attention of the subject. Because active attention by the subject is required, only one such sequence at a time may be controlled without interference, unless two sequences each requires such a slow sequence of activations that they can be interwoven serially. Controlled processes are therefore tightly capacity-limited, but the cost of capacity limitations is balanced by the benefits deriving from the ease with which such processes may be set up, altered, and applied in novel situations for which automatic se-

[6] It may also be possible to lower the threshold by a recent activation of the same automatic sequence. However, this hypothesis is difficult to distinguish from the contextual one, because an activation of a sequence is usually accomplished by the prior, recent presentation of information associatively related to the sequence, and this related information is still likely to be in STS at the time of test.

quences have never been learned. Controlled processing operations utilize short-term store, so the nature of their limitations is determined at least in part by the capacity limitations of STS.

6. *Short-Term Store (STS)*

Short-term store is the labile form of the memory system and consists of the set of concurrently activated nodes in memory. The phenomenological feeling of consciousness may lie in a subset of STS, particularly in the subset that is attended to and given controlled processing.[7]

The capacity of STS is determined stochastically (see Shiffrin & Cook, Note 3) so that a large amount of information may be present (activated) at any one moment, but only a small amount of information will persist for an appreciable time period of several seconds or more. Forgetting or loss from STS is simply the reversion of currently active information to an inactive state in LTS.

What determines the loss rate? We suppose that the rate of loss of any informational element or node in STS depends on the number of similar elements simultaneously active in STS. By similarity we refer not only to formal physical similarity but also to similarity of features at comparable levels of processing (e.g., the typeface of a printed word will be less likely than the word itself to cause forgetting of a verbally encoded second word in memory). When a large amount of similar information is present, the loss rate will be rapid but will slow as the amount of active information decreases. To give an example, when a complicated visual scene is presented to a subject briefly in a tachistoscope, a flood of visual information enters STS and initiates a series of chains of automatic processing that result in higher level features' also entering STS. However, most of the activated information will have decayed and will be lost from STS in just a few hundred milliseconds after physical offset of the scene (see Sperling 1960); just a few of the features, perhaps at higher levels, will remain present longer than a few seconds.

STS has two somewhat distinct roles. The first is the provision of a temporary storehouse for information currently important to the organism. That is, it acts as a selective window on LTS to reduce the amount of information for processing to manageable proportions. The second role of STS is the provision of a work space for decision making, thinking, and control processes in general.

7. *Learning: Transfer to LTS*

Consider first what is meant by transfer from STS to LTS. Transfer implies the formation in permanent memory of information not previously there. To be precise, this consists of the association (in a new relationship) of information structures already in LTS. A minimal requirement for this new associative structure is the simultaneous presence in STS of the separate elements to be associated or related. That is, the various nodes to be linked in a new relationship must be activated in STS. Thus, transfer to LTS does not imply the removal of the transferred information from STS, nor the placing of new "subunits" in LTS that do not already exist in either LTS or STS. Rather, transfer means the formation of new associations (or the strengthening of old associations) between information structures or nodes not previously associated (or strengthened) in LTS. Most new associative structures will include as a component the context in STS at the time of the transfer.

The above remarks specify the nature of new learning in STS but not the cause underlying the storage mechanism. It has been

---

[7] An alternative formulation, closer to that suggested by Atkinson and Shiffrin (1968), would use the term *STS* to refer to those nodes that are given attention and controlled processing. Other activated nodes would be referred to by another term, such as "sensory register" in Atkinson and Shiffrin's terminology. At such a general level of discussion, it is doubtful that there are substantive differences between the two approaches—versions of each could be constructed to be identical to each other. Each approach has its own heuristic advantages, but a theory is better judged in terms of its detailed assumptions and predictions than by its choice of either type of terminology.

shown that rehearsal (and coding rehearsal) are strongly implicated in learning (see Shiffrin, 1975b, for a discussion). We prefer an extension of the rehearsal hypothesis. We assume that what is stored is what is attended to and given controlled processing. Rote rehearsal is just one form that attention can take. In fact, maintenance rehearsal may result in storage of low-level auditory or verbal codes not helpful for long-term recall, (but helpful for long-term recognition), while coding rehearsal may result in storage of high-level codes useful for recall. This model suggests that some degree of attention or controlled processing is a prerequisite for storage. Thus, incidental learning situations, in which no extended rehearsal takes place, will still result in some learning to the degree that the input items are attended to during presentation.

These hypotheses that controlled processing will underlie learning apply of course to the development of automatic processing, and in particular, to the development of automatic detection studied in the search and detection paradigms earlier in this paper and in Part I. In those paradigms it was seen that automatic detection developed only with consistent training. Consistent training is crucial when the learned sequences in LTS contains an internal or overt response that will be harmful to performance on trials that are inconsistent. For example, an attention response to an item will harm performance on trials when that item is a distractor. Purely informational sequences (i.e., sequences that do not include responses that direct internal processes or overt actions) will be stored in LTS when attended to, regardless of consistency of training. In all cases, however, consistent training and large numbers of repetitions should lead to stronger automatic encodings and processes.

An implication of the hypothesis that transfer to LTS is engendered by controlled processing is the important concept that controlled processing, and hence attentional limitations, will be involved during the acquisition of automatic processing. We have largely been identifying attentional limitations with controlled processing (e.g., controlled search) and have shown how automatic processing (e.g., automatic detection) can bypass attentional limitations. It should not be overlooked that the initial learning of automatic processing may require a variety of control processes and hence will be subject to various limitations that may disappear when learning has progressed to a high level.[8]

8. *Retrieval of Information from Short-Term Store*

Several retrieval modes from STS are possible. An automatic process might direct the retrieval process to just a subset of the activated information (e.g., only the letters, not the masks, are compared in our VM search studies).

The various controlled search strategies that are available are normally serial in nature but a variety of search orders is possible. Search order may depend upon instructions, strength of short-term traces, the nature of categories of the traces, the modality of the information, and the structure of STS. This last point is worth emphasizing: Since STS is embedded in LTS, it partakes of the structure of LTS and is not a totally undifferentiated mass of information. Thus the ordering of a search can utilize this existent structure. In addition, the comparison process may be exhaustive or terminating, and the information located in one phase of the search can be used to redirect a later phase of the search.

It is the control of search order that is responsible for many of the selective attention effects that are observed. Information in locations to which attention is first directed (i.e., the first locations searched) will in general receive faster and more accurate processing for the obvious reasons.

[8] This distinction is helpful in understanding the relation of our work to studies of attention in infrahuman organisms: the traditional studies of attention in discrimination learning and blocking (e.g. see Mackintosh, 1975) are involved with those limitations that occur during acquisition, while a number of newer studies are involved with limitations in steady state situations when learning is not possible (e.g. see Riley and Leith, 1976).

9. *Retrieval of Information from Long-Term Store*

Both automatic and controlled processing are utilized in long-term retrieval. In general, an automatic associative mechanism will cause currently inactive information in LTS to become active when associatively related information is in STS. The simplest example is the perceptual encoding process by which information concerning environmental inputs is located in LTS. As each feature is activated, it can serve as a base for further activation of features associatively related to it and to other previously activated features. This same principle applies to long-term retrieval of higher order information.

It should be noted well that information once activated automatically from LTS will then be in STS and may have to be located through use of any of the short-term retrieval mechanisms discussed above.

How does the subject exert control over LTS retrieval? The subject can utilize rehearsal and selective processes to raise the strength of certain information in STS above the strength of other information. Then the activated and retrieved information from LTS will tend to be information related to the selectively accentuated subset.

In Shiffrin's (1970) paper the details of the controlled LTS search were presented at length. In summary, the controlled search process is a series of cycles. On each cycle (a) the subject generates probe information and places it in STS; (b) the probe information, along with the general contextual information presently in STS, activates associated information from LTS, called the *search set;* (c) the subject searches the STS search set; (d) the subject decides whether the appropriate information has been found and whether the search is over. The process then continues cycling. Note that step (c) is a retrieval from STS, so that one phase of controlled LTS retrieval may be STS retrieval.

10. *Long-Term Forgetting*

The forgetting of information in LTS is, by definition in the present theory, a question of retrieval failure. The retrieval processes discussed in the preceding sections are by their nature imprecise and fallible. When retrieval fails, perhaps temporarily, then we say forgetting has occurred. The causes of such failure are somewhat to the side of the main interests of this paper but may be summarized briefly as follows. First, the probe cues used to activate related information in LTS may be ineffective, either because the subject chooses them incorrectly or because the general context in STS at the moment is dissimilar to that stored with the desired information in LTS. Second, the process of retrieval itself may harm additional retrieval attempts, especially if the probe cues are not altered from one cycle of the search to the next. Third, the search may fail due to premature termination; that is, it may fail due to a decision that further retrievals are not worth the effort, or due to the reaching of some sort of time limit available for search. (See Shiffrin, 1970, 1976, for a fuller discussion of these matters.)

B. *The Characteristics of Controlled and Automatic Processing*

In capsule form, we have covered the major phases of the information-processing system. We shall now attempt to define and outline the characteristics of automatic and controlled processing in a more precise fashion.

1. *Controlled Processing*

It is important to note that not all control processes are available to conscious perception, and not all control processes can be manipulated through verbal instruction. It is therefore convenient to divide control processes into two classes: *accessible* and *veiled.* Accessible control processes are those like rote rehearsal or alphabetic search of LTS that can be instituted and modified by instruction. These are generally slow processes that are easily perceived by the subject. Veiled control processes are those like the serial comparison of items in short-term store that are difficult to modify through instruction. They are not easy to perceive through introspection because they take place so quickly.

Both classes of control processes have the following general properties: (a) They are limited-capacity processes requiring attention. Because these limitations prevent multiple control processes from occurring simultaneously, these processes often consist of the stringing together in time of a series of singly controlled unitary operations. (b) The limitations of control processes are based on those of STS (such as the limited-comparison rate and the limited amount of information that can be maintained without loss). (c) Control processes can be adopted quickly, without extensive training, and modified fairly easily (though not always by verbal instruction). (d) Control processes can be used to control the flow of information within and between levels, and between STS and LTS. In particular, they can be utilized to cause permanent learning (i.e., transfer to LTS) both of automatic sequences and of information in general. (e) Control processes show a rapid development of asymptotic performance. For a given control process, if automatic processing does not develop (due, say, to varied-mapping procedures) and if the constituent elements of the process do not otherwise change due to long-term learning, then performance level will stabilize very quickly at an asymptotic value. When performance improves over trials, it does so because the control process is changed, or because automatic processing develops, or because the constituent nodes that are linked by the control process are themselves altered by long-term learning.

Common examples of control processes include maintenance or rote rehearsal, coding rehearsal, serial search, long-term memory search, and decisions and strategies of all kinds. These will be discussed in more detail in the sections below.

2. Automatic Processing

Automatic processes have the following properties. (a) They are not hindered by the capacity limitations of STS and do not require attention. Thus automatic processes often appear to act in parallel with one another and sometimes appear to be independent of each other. (b) Some automatic processes may be initiated under subject control, but once initiated all automatic processes run to completion automatically (though some indirect control is possible through manipulation of the contents of STS at the time of the inciting input). (c) They require considerable training to develop and are most difficult to modify, once learned. (d) Their speed and automaticity will usually keep their constituent elements hidden from conscious perception. (e) They do not directly cause new learning in LTS (though they can indirectly affect learning through forced allocation of controlled processing). (f) Performance levels will gradually improve over trials as the automatic sequence is learned. In many cases asymptotic performance levels may not be reached for thousands of trials.

It is particularly important to specify carefully what we mean when we say that an automatic process does not partake of the capacity limitations of STS. To make this clear, let us suppose that X, Y, and Z are nodes that are automatically activated in turn by input I in the presence of general context nodes C. We may depict the sequence as follows:

$$\begin{tikzcd} I \arrow[r] & X \arrow[r] & Y \arrow[r] & Z. \\ C \arrow[ur] & C \arrow[ur] & C \arrow[ur] & \end{tikzcd}$$

As each of X, Y, and Z is activated, it enters STS and its future residence in STS will be affected by the limitations of STS. In fact, if the nodes X, Y, and Z do not include an attention-attracting response, then it is quite conceivable that all three nodes will decay and be lost within a few hundred milliseconds, and the subject will be quite unaware that such an automatic sequence ever occurred. Nevertheless, the sequence itself is not governed by the limitations of STS in the sense that the probability of its being activated and the rate of its occurrence will not be affected by other concurrent automatic and conrolled processes taking place in STS (at least if context C is present and the other concurrent processes do not also utilize nodes X, Y, or Z).

3. The Development of Automatic Processing

The tendency for one node to activate another will be increased if both nodes are present simultaneously in STS and if a control process and attention are directed toward these nodes, thereby increasing their salience in STS.

Before we can describe efficacious training conditions, we must distinguish between two types of automatic processes. One type of automatic sequence may be called *actional,* because it includes phases that direct internal processes (like calls for attention) or phases that produce overt responses (such as button presses). A second type, called *informational*, contains no directions for actions. The distinction is crucial because an informational sequence strengthened on one trial will not have deleterious consequences for performance on other types of trials. On the other hand an actional sequence may give rise to a response antagonistic to a response required on another trial. Thus, to be useful in a given task, actional sequences require special, consistent, training conditions.

To make this point clear, suppose that node B produces a response antagonistic to that produced by node C, while nodes D and E produce no responses. If any of A–B, A–C, A–D, or A–E is trained alone, then that sequence can be learned. If training on A–D is mixed from trial to trial with training on A–E, then A will come to elicit *both* D and E. However, if A–B trials are mixed with A–C trials then learning of both will be impossible, since A cannot lead simultaneously to two antagonistic responses. The automatic detection system discussed at length in this paper is just an actional sequence, since an automatic-attention response to one stimulus is incompatible with a simultaneous attention response to another stimulus; therefore automatic detection requires consistent training to develop. Thus, in the varied-mapping conditions of any of our studies automatic detection could not develop and controlled processing had to be utilized.

4. *Combinations of Automatic and Controlled Processing*

Although sensory inputs are first encoded with the automatic processing system, with the results being made available to controlled processing, it must not be concluded that automatic processing invariably precedes controlled processing. In fact, automatic and controlled processes can proceed in parallel with one another (as in Experiment 4d when automatic processing of the elements on the invalid diagonal proceeded in parallel with controlled processing of the valid diagonal). Even more important, controlled processing is often used to initiate automatic processing. Particularly in complex processing situations, (such as reading), an ongoing mixture of controlled and automatic processing is utilized. The next steps in the serial, controlled processing sequence are based on the output of automatic processes initiated earlier; then new automatic sequences are initiated and these run to completion in parallel with the ongoing controlled processing.

5. *The Benefits of Automatic and Controlled Processing*

A system based on the two basic processing modes with the characteristics we have described has many advantages. In novel situations or in situations requiring moment-to-moment decisions, controlled processing may be adopted and used to perform accurately, though slowly. Then as the situations become familiar, always requiring the same sequence of processing operations, automatic processing will develop, attention demands will be eased, other controlled operations can be carried out in parallel with the automatic processing, and performance will improve. Some of the advantages of such a system are (a) It allows the organism to make efficient use of a limited-capacity processing system. The development of automatic processing allows the limited-capacity system to be cleared and devoted to other types of processing necessary for new tasks. (b) It allows attention to be directed (through automatic-attention responses) to important stimulation, whatever the nature of the ongoing controlled processing. (c) It allows the organism to adjust to changes in the environment that make previously learned activity patterns useless or harmful. (d) It allows the organism to deal with novel situations for which auto-

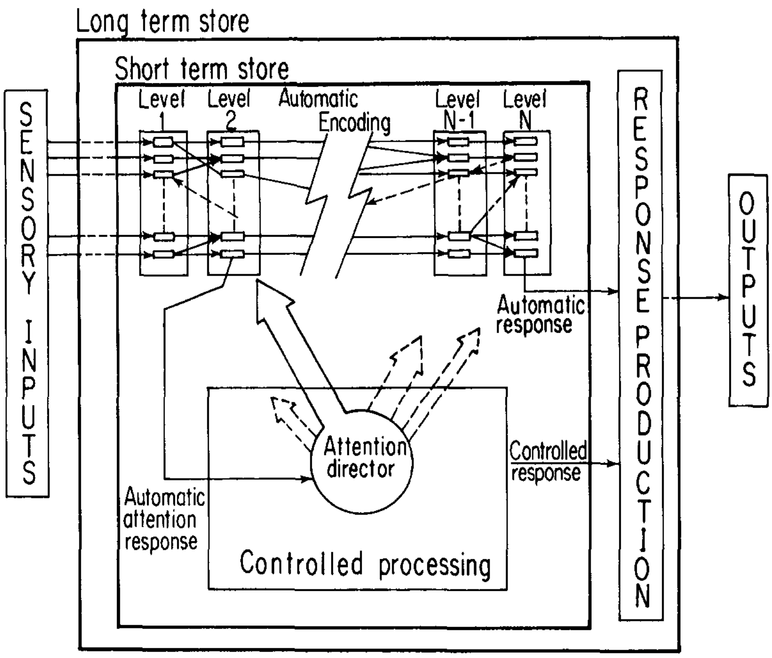

Figure 11. A model for automatic and controlled processing during tasks requiring detection of certain input stimuli. Short-term store is the activated subset of long-term store. N levels of automatic encoding are shown, the activated nodes being depicted within each level. The dashed arrows going from higher to lower levels indicate the possibility that higher level features can sometimes influence the automatic processing of lower level features. The solid arrow from a node in Level 2 to the attention system indicates that this node has produced an automatic-attention response, and the large arrow from the attention system to Level 2 indicates that the attention system has responded. The arrow from level N to the Response Production indicates that this node has called for an automatic overt response, which will shortly be executed. The arrow from Controlled Processing to the Response Production indicates the normal mode of responding in which the response is based on controlled comparisons and decisions. Were it not for the automatic responses indicated, detection would have proceeded in a serial, controlled search of nodes and levels in an order chosen by the subject.

matic sequences have not previously been learned. (e) It allows the organism to learn increasingly complex modes of processing by building upon automatically learned subsystems.

C. *A Framework for Processing in Detection, Search, and Attention Tasks*

The framework we have in mind is built within the general theory already presented. It is illustrated in Figures 11 and 12. When a set of inputs is presented (let us suppose for convenience that the inputs are visual) then each begins to undergo automatic processing. The system automatically encodes each stimulus input in a series of stages and activates a series of features in the process. For example, the letter "M" may first be encoded in features indicating contrast, color, and position; then curvature, convexity, and angles; then a visual letter code and a verbal, acoustic-articulatory code, then the codes "letter," "consonant," "capital;" and finally, perhaps, semantic and conceptual codes like "followed by 'N'," "middle of the alphabet," and the like. (We do not necessarily imply that these features are correct or exhaustive; if, say, amplitude components of the spatial-frequency analysis of the inputs prove to be relevant features encoded by the system, the theory we describe would be unchanged in all important respects.) What features will become activated depends on the physical nature of the nervous system that was predetermined genetically, the degree and type of prior learning, the physical characteristics of the display (like duration and contrast), and the general context, both that in the environment and that generated in STS by the subject. To the extent that the subject directs the sensory receptor orientation and to the extent that internally generated information can alter the context in STS, the subject will have at least some indirect control over automatic sensory coding.

The automatic processing as described above takes place in parallel for each of the input stimuli. The processing of each stimulus is often independent, except for lateral and temporal interactions at early stages, called masking, and except for learned relationships between items that may affect processing at later stages (if adjacent letters form a word, for example). The various features that are activated are all placed thereby in STS where they reside for a short period before being lost (i.e., before returning to an inactive state in LTS).

We propose that some of the features or nodes that are automatically activated may initiate a response that will direct subsequent processing or subsequent actions. For example, an attention response might be activated by a feature; the attention response might direct controlled processing to the corresponding set of features representing that input stimulus, so long as other competing attention responses do not occur simultaneously. As another example, an overt response might be engendered by a particular stimulus (such as a startle response to a sudden loud noise, a galvanic skin response to an aversively conditioned word, a ducking response to a missile thrown at the head). If the set of inputs contains a target stimulus that gives rise to a relevant nonconflicting attention response, or to an overt response, then we say that automatic detection is operating. Note that the attention response could be attached to features at any level of processing and to more than one feature at a time. In Figure 12, as an example, attention responses are attached both to the feature "8" and to the feature "number."

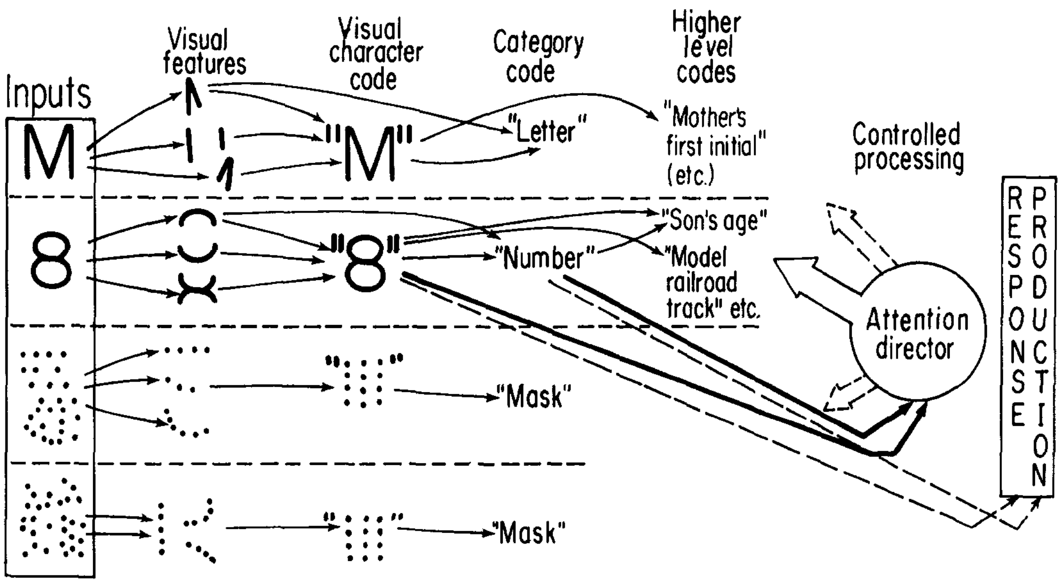

Figure 12. Another view of the model shown in Figure 11. This figure depicts a conceivable (but abbreviated) series of feature encodings for a frame in which two characters and two masks are presented. The consistent-mapping condition is shown in which numbers are memory-set items and letters are distractors. The arrows skipping levels indicate that a given feature can help activate features at several different levels. The stimulus *8* is a consistent-mapping target, and hence both the visual and category codes produce automatic-attention responses. The attention response has attracted attention to the information deriving from the input *8*.

It is possible that many of the inputs will give rise to attention or other responses and that these responses will conflict. For example, every input might automatically engender an attention response; since it is impossible to direct attention to every item in the display at once, the various attention responses will conflict and cancel each other. In such cases, the subject will not be governed by the output of the automatic system and will base subsequent actions and responses on the results of a controlled search. These hypotheses were supported in Experiment 2 (see Figure 4, left panel for the results). In this study, *all* inputs were items that subject had previously been trained to respond to with automatic-attention responses. The results showed that the subjects reverted to the use of controlled search; furthermore, the controlled search was almost identical to that seen in the usual VM conditions in which none of the inputs had been trained to elicit automatic-attention responses.

In cases when the input stimuli do not activate automatic responses that govern the detection process, then a controlled search must be used. In this event, the subject must attempt to search through the set of features that is activated and to use a search strategy that is as efficient as possible. For example, in our visual search studies, letter codes are normally compared in preference to (and prior to) sets of angles and lines; a category code representing a set of letters is often compared in preference to (and prior to) individual letter codes, as shown in Experiment 3. Furthermore, the order of search and the placement of decisions within the search are also controlled in an attempt to increase efficiency. To give an example, comparisons in our studies in Part I took place in an order that cycled through the frame for a given memory-set item before switching to a new memory-set item, and a matching decision was made after every comparison. To give another example, matching decisions in many memory search studies (M varying, $F = 1$) are withheld until all comparisons are completed (Sternberg, 1975).

Note finally that many attention tasks utilize search or detection paradigms. Whenever this is the case, the framework described here applies equally to attention tasks. Divided-attention tasks study the increments in performance that may occur when the load is decreased. Decreased loads are usually specified through advance instructions that certain inputs are irrelevant, thereby decreasing the size of the memory set, the frame set, or both.

According to our framework, experiments that do not show dividing attention to reduce performance fall into two classes. One class includes those studies in which automatic detection is operating. These are usually studies utilizing a CM paradigm and high degrees of practice. Divided-attention deficits will not be seen because the target stimulus will be detected automatically, in parallel with the other stimuli, and often independently of other stimuli. An exception to this general rule may apply when two (or more) targets are presented simultaneously on a trial. Even if both targets generate attention responses, there may be a difficulty in discriminating a double from a single occurrence. In such an event, the controlled comparison system may have to be called into play to count the targets. (A detection decrement may then occur, because the second item may have decayed from STS by the time the comparison of the first item is complete; see Moray, 1975, Shiffrin, 1976, and the discussion of the 0 spacing effect in the CM condition of Experiment 3, Part I.)

The second class of studies in which divided-attention deficits do not occur is that in which controlled search is utilized but the capacity of STS is not stressed. Examples are seen in the series of studies by Shiffrin and his colleagues, summarized in Shiffrin's (1975a) article. In such studies either the load in STS is kept low, or a cue informing the subject which input is relevant appears at about the same time as the inputs.

It is, however, only in specially designed situations that capacity limits are not stressed. In most studies requiring controlled search, the extra time required to search a larger number of relevant inputs will impair performance, whether measured by reaction time or accuracy. This was the case in the

VM conditions of the experiments reported in this paper, and in Part I: In all these cases increases in load reduced performance. There are two basic reasons for the performance reduction in controlled search studies. First, some of the features in STS may decay and revert to LTS before the comparison process reaches them. Second, new inputs may arrive and require processing, forcing the comparison process to switch away from the previous inputs. Probably only this second factor operated in the studies we have reported in this paper. For either of these reasons, divided-attention deficits can be expected in detection tasks requiring controlled search.

The situation is quite different, however, in focused-attention studies. In these studies, certain inputs are known to be relevant, and others are to be ignored. If controlled search is being utilized, there is little reason to expect that any substantial deficit will be caused by the presence of to-be-ignored stimuli, at least if it is clear to the subject well in advance which stimuli and locations are relevant and if there are no location confusions. In such cases the subjects should direct the search order so that the relevant locations are searched first; thus performance deficits should not arise. Evidence supporting these propositions is found in Experiment 4a: Subjects were able to search one diagonal and ignore the other.

The tasks in which large deficits in focusing attention will occur are those in which irrelevant items give rise to attention responses during automatic encoding. In such cases, attention will be attracted to the to-be-ignored item and a loss of time will occur before the controlled search can be redirected to the relevant inputs. The time lost will impair performance. Such focused-attention deficits were seen in the reversal results of Experiments 1 and 2, and in the attention results of Experiment 4.

Thus, our framework can be applied to a wide variety of search, detection, and attention tasks, although additional assumptions must be made to generate precise models to deal with the particulars of each task. In Part I we presented a quantitative serial, terminating search model and discussed some alternative models. Our discussion was limited, however, to our basic single or multiple-frame search task with small values of M and F. We will now consider a few alternative paradigms, along with special assumptions they might require.

1. *Tasks Utilizing Large Memory Sets*

Suppose there are too many items in the memory set to be maintained in STS. Of course, then, the memory set must be learned in LTS in advance of the test. There are then two basic search modes. In the first, the test item is automatically encoded and accesses a node in memory that contains the desired information. For example, in deciding whether an item is any word, the input item, if a word, is encoded to the nodes representing the word and the information in those nodes is usually sufficient to classify the item as a word; on the other hand, a nonword obeying appropriate orthographic constraints is encoded to a lesser degree and the failure of the automatic encoding process could itself serve as a cue for "nonwordness." The second type of search mode would consist of entering successive parts of the memory set from LTS into STS, using controlled search to examine each subset in STS. For example, if asked whether there is a flower, country, or first name whose fourth letter is "u" a subject might successively generate members of each category and serially check them.

The above comments make it clear that the case of large memory sets is not different in principle from the cases in which the memory set may be held in STS, though the models may need additional assumptions. For example, Atkinson and Juola (1974) carried out a study in which subjects had a large memorized set in LTS. They proposed a hybrid model in which a judgment based on familiarity enabled the subject to decide on some trials without a search that the test item was definitely in, or definitely not in, the memorized set. On trials when this initial familiarity judgment was ambiguous, then a controlled search through the set was made.

2. *Large Display Sets*

When large numbers of items appear in the display, then it may turn out that acuity varies considerably across the display for a given fixation point. In turn, acuity changes are likely to lead the subject to adopt a strategy in which eye movements are used to bring subsets of the display successively into a high-acuity region. The location of each new fixation can be governed by a controlled search strategy or by an automatic process, but in either event the necessity of obtaining good acuity will force this process to be serial across successive eye movements. Within each fixation, furthermore, either automatic detection or controlled search could be utilized. The search studies in which subjects scan rows of characters for targets have shown both modes of search. For example, the Neisser (1963) results, reviewed in Part I, showed that automatic detection developed for subjects well trained in responding to letter sets presented in CM fashion (memory-set size had no effect), while controlled search was utilized for sets that had been given only small amounts of training. An interesting intermediate case occurred when the subject was instructed to locate a row of characters that did *not* contain a given character. In this case, even with CM training, the task proved quite difficult. In our view the subject would have to adopt a serial process from row to row (rather than from fixation to fixation). Within each row automatic detection might be used to locate the key character, but then a decision would have to be made before the next row could be considered.

3. *Categorical Partitioning of Display Sets*

Partitioning of the display set has usually been accomplished by a categorization dependent on a relatively gross simple physical feature, such as color, shape, or size. When displays are segregated into two or more groups by such features, processing may be affected in two ways. First, the patterning of the input can govern the nature of automatic processing—controlled processing may be directed automatically to a subset of the inputs (assuming that one particular subset is consistently relevant). Second, the perceptual categorization could influence the order and perhaps the nature of controlled search. To give one simple example, a display arranged in two separate rows will tend to be processed one row at a time, in reading order. A second example occurs in situations where controlled processing can enable one to more quickly decide about the category of the input than about the memory-set membership of the input. In these cases a two-phase controlled search might be adopted, with the first phase locating the relevant inputs by category, and the second phase matching these inputs to the memory set. This is illustrated by Green and Anderson's (1956) and Smith's (1962) studies in which a two-digit number of a particular color had to be found in a display of many two-digit numbers of differing colors.

Automatic processing based on perceptual categorization may have played an important role in the studies of search and detection reported in this paper. We have assumed throughout that only the display positions with characters, and not those with masks, are considered in the search. How is search restricted to character positions? It is conceivable that a preliminary controlled search is used to identify character positions, but the results of Experiment 4a argue against such a view—knowing the relevant diagonal of a four-character frame leads to performance identical to that seen when two masks and two characters appear randomly on each frame. Thus, a preliminary controlled search would have to be extremely fast relative to the time for each character comparison. More likely, an automatic process develops that directs the controlled search away from mask positions and toward character positions. Since such a response would be consistently trained and reinforced, it should be learned in a fashion similar to that for automatic-attention responses.

It should be noted that the development of automatic responses to direct controlled search is very similar to the process termed by Neisser (1967) as "pre-attentive." Neisser was trying to explain how search could be directed to a perceptually segregated subset of the inputs. We are suggesting that an automatic process might develop and direct the search in situations where the search-directing response is consistently trained (as is usually true in such studies). We do not feel, however, that the segregation of the relevant inputs must be based on simple physical features. With enough consistent training, an automatic search-directing response could develop for a set of stimuli defined by quite complex features. For example, Gould and Carn (1973) reported results suggesting that an automatic process was directing search to the locations of items from a set of 10 potentially relevant letters and away from the locations of items from a set of 8 other letters that *never* included a target. (Any target was always drawn from the potentially relevant set of 10 letters, but on most trials distractors were chosen from this set of 10 as well as from the set of 8).

Although the evidence is not yet available, we raise the possibility that the converse of the argument in the preceding paragraph might also hold: Even with perceptual segregation of the display into categories determined by simple physical features, automatic restriction of the search to the relevant category might not be possible unless the relevant category is consistent across trials.

4. *Attention Tasks*

The models applying in attention studies that utilize detection or search paradigms are just those discussed in the preceding sections. The general rule is that instructional manipulations affect the order of controlled search; increases in load cause the time needed for controlled search to increase; and consistent training leads to the development of automatic responses that allow attentional limitations to be bypassed.

We will say a few words, however, concerning attention studies that do not utilize detection or search paradigms. One class of studies utilizes a memory paradigm. The inputs are manipulated in ways similar to those used in detection studies, but the subject is asked to recall or recognize the items at some time after presentation. One example is the "split-span" technique reviewed by Broadbent (1958, 1971). Other examples are reported by Sperling (1960) and Von Wright (1970), who studied cued partial report following tachistoscopic exposure of character displays.

In these memory paradigms, essentially the same underlying mechanisms apply as in the detection paradigms, though memory storage rather than detection is the processing goal. When controlled processing is utilized, the subject has considerable control over storage: The items that are rehearsed or coded are the items that will be recalled. If an input is presented that generates an automatic-attention response, then this input will receive attention and tend to be recalled regardless of the nature of the controlled process utilized to store the other items. An excellent demonstration of these points may be found in Kahneman (1975).

5. *Threshold Detection Tasks*

An input in a threshold detection task, by definition, is presented in such a way that automatic encoding on most trials will be incomplete and therefore ambiguous. That is, in Section IV.A.4 we discussed the fact the input stimulation must be above some threshold level for automatic encoding to run to completion in an accurate fashion. In Figure 12, for example, the "M," if presented near threshold, might cause activation of some of the features at the visual feature level, but no features at higher levels. This set of partial features will in general be ambiguous —several letters and numbers might be consistent with the activated feature set. Thus, a target will on some trials give rise to feature sets consistent with either targets or distractors. Similarly, distractors will on some trials give rise to feature sets consistent with either targets or distractors. Under these conditions it is clear even in CM situations that automatic-attention responses and automatic detection can only develop very slowly, if at all, because the internal representations of targets and distractors will be mapped consistently to responses only on those few trials when the encoding happens to be complete.

A similar argument applies if encoding is inaccurate rather than incomplete.

Even in CM threshold tasks, therefore, the subject will be forced to adopt controlled search on most of the trials, namely the trials on which encoding is incomplete. On such trials a serial comparison process will be necessary to match the activated features against the features of the members of the memory set. Occasionally, however, automatic encoding will manage to process the target completely. Suppose that almost always, those stimuli that are completely processed will be encoded correctly. Suppose also that a CM procedure is being utilized in the task. In these special conditions, an attention response can be learned that might be utilized on those trials when complete encoding of the target happens to occur.

A model very similar to this was used by Shiffrin and Geisler (1973) to fit the results from a threshold detection study by Estes (1972). In that model, stimuli automatically processed to a final, incomplete result were called "component detections" and stimuli automatically processed to the complete (letter) level were called "letter recognitions." The model proposed a two-part controlled search in which the letter recognitions were searched first (in parallel through memory, but serially through the display), and then the component detections were matched serially, feature by feature, against the members of the memory set. In light of the present work, we would like to propose for consideration an alternative but similar model in which the letter recognitions result in an automatic-attention response, so that controlled search need be utilized only if the target is not among the letter recognitions. In practice, the predictions of these two models would not differ greatly.

In summary, it seems clear that controlled search, whether at the feature or character level, is bound to be the predominant detection mechanism in all threshold detection tasks, but there is a possibility in CM paradigms that automatic detection can be utilized on a subset of the trials when the target happens to be encoded relatively completely.

Our general view of processing, in threshold detection tasks as well as many others, can be summarized briefly in the following manner: Sensory inputs are encoded automatically, resulting in states of evidence (consisting of feature sets) that the subject utilizes in subsequent controlled processing. While generally accurate, this view is too simple in several respects. First, it must not be assumed that there is ever one moment in time at which all activated features are simultaneously present and available to the subject for controlled processing. Rather it will usually be the case that some features will still be undergoing the process of encoding (and hence will not yet be activated) while other features have already been activated and are in the process of being forgotten. Second, it must not be assumed that controlled processing will follow automatic processing in time. Such an assumption is a fairly accurate simplification when used in models for certain tasks including threshold detection. However, it is an inherent feature of our theory that automatic and controlled processing can often operate in parallel, so that controlled processing can begin while automatic processing is still progressing. This point becomes crucial in situations in which the subject is required to give extremely fast responses (including those paradigms presenting above-threshold stimuli), since a very fast response might have to be based on the features available at a certain point in time, even though better features might be available at a later point in time.

D. *Automatic and Controlled Processing in Memory Storage and Retrieval*

In this section we shall review briefly how automatic and controlled processing might be involved in short-term retention, learning, and in long-term retrieval.

1. *Maintenance of Information in Active Memory*

We suggest that the control process most evident to introspection is rehearsal. Atkinson and Shiffrin (1968) studied many tasks in which subjects utilized a continually updated rehearsal "buffer" to aid their memory system. Rote rehearsal is of course only one of many control processes that cause items to have an extended residence in STS; coding, deciding, retrieving, and the like also cause a similar result, though these processes are primarily intended to serve other purposes. In general, items given attention of any sort are maintained in STS, including items to which attention is drawn by an automatic process.

Items from which attention has been removed do not necessarily leave STS quickly, however. When the load is small and new inputs are minimized, items can remain in STS for extended periods (up to 40 sec in the Shiffrin, 1973, study). Another example is seen in the overt forced-rehearsal study reported in Atkinson and Shiffrin (1971). A long list of items was presented and subjects rehearsed aloud a continually updated list of the three most recent items. As a result, the last three items presented were always recalled on an immediate test. However, the probability of recall for items prior to the last three decreased systematically as a function of the spacing from last rehearsal to test. Thus an item removed from controlled processing still requires a period of time before it becomes lost from STS (i.e., becomes inactive). This persistence in the absence of attention might be called the automatic component of short-term maintenance. It is this automatic component of persistence that was studied by Shiffrin and Cook (see Note 3) and that determines the basic capacity and persistence of STS. In brief, Shiffrin and Cook propose that the loss process is a stochastic mechanism whose rate increases as the amount of similar material concurrently in STS increases.

2. *Coding, or Transfer to LTS*

The role of control processes in long-term storage is well established. Rote rehearsal appears to transfer low-level auditory-verbal codes to LTS; these are not very useful for long-term recall but can facilitate long-term recognition. Coding rehearsal appears to facilitate long-term retrieval of all kinds. In general, what is attended to and rehearsed is what is stored in LTS. (See Bjork, 1975; Craik & Jacoby, 1975; Craik & Lockhart, 1972; and Craik & Tulving, 1975, for a discussion of these issues.)

The other side of the storage question concerns the nature of LTS transfer when attention is not directed to an item or sequence of items. Studies of incidental learning tend to show that subjects learn what they attend to, not what they are told to learn (see Craik & Tulving, 1975; Hyde & Jenkins, 1969, 1973; Postman, 1964; Schneider & Kintz, 1967).

If controlled processing is necessary for long-term learning then automatic processing without controlled processing should not lead to appreciable retention. One source of evidence is found by examining retention for distractors in CM tasks. Such items presumably have received almost no controlled processing. However, it may be assumed that due to prior exposures and prior learning the distractors are given automatic encoding at least up to the "name" level (for evidence see Corteen & Wood, 1972; Keele, 1972; Von Wright et al., 1975, among others). Does the automatic encoding that these distractors receive during the many trials of a CM task lead to any retention? Moray (1959) had subjects repeat aloud a prose passage in one ear while a seven-word list was repeated 35 times in the other ear. Later recognition for the unattended words was at the chance level. Gordon (1968) showed that a set of four distractors in a CM search task was recognized near the chance level even after 10 days of practice. Gleitman and Jonides (1976) showed that distractors were more poorly recognized in a CM search task than in a VM search task (though, due to the low levels of practice used in their study, the effect could have been due to the use of a controlled search for categories in the VM condition).

Evidence of a rather different sort is found in reading tasks. If automatic processing does not lead to retention, then it might be expected that a skilled reader who automatically processes surface-structure features and who gives controlled processing to conceptual properties of a passage would retain little information regarding the surface-structure details. Bransford and Franks (1971) have collected data that may be interpreted in this fashion. Of course, if controlled processing is directed toward levels of analysis normally carried out automatically, then the features attended to will be remembered (see Postman & Senders, 1946).

To summarize, whether or not some nominal storage manages to be eked out in the absence of controlled or attentive processing, it seems clear that any phenomenon that would be appropriately called "automatic storage" is a far less important determinant of LTS learning than controlled processing.

Before leaving the general topic of learning it is interesting to speculate on the development of complex information-processing skills. Since controlled processing is limited, only a small part of the memory system can be modified at any one time. After invariant relations are learned at one stage, the processing becomes automatic and controlled processing can be allocated to higher levels of processing.

For example, the child learning to read would first give control processing and then give automatic processing to various units of information. The sequence of automatically learned units might be foreground–background, features, shapes, letters, words, and meanings of phrases or sentences. The child would be utilizing controlled processing to lay down "stepping stones" of automatic processing as he moves on to more and more difficult levels of learning. The transition from controlled to automatic processing at each stage would result in reduced discrimination time, more attention to higher order features, and ignoring of irrelevant information. Gibson (1969, chap. 20) describes these effects to be three of the major trends in perceptual development. In short, the staged development of skilled automatic performance can be interpreted as a sequence of transitions from controlled to automatic processing.

3. *Retrieval from LTS*

Retrieval from LTS has a large automatic component. Sensory inputs result in an automatic series of stages of encoding that activate many features (and perhaps responses) and place them in STS. Internally generated inputs also tend automatically to activate associated information in LTS and to place it in STS.

In general, the two primary controlled phases of LTS retrieval consist of the processes the subject uses to search and process the information in STS that has just been activated, and the processes used to alter the probe information from step to step of the LTS search.

It is important to note that the selection of probe information is not entirely under subject control. In fact, the probe information includes the general contextual information present in STS at the time of retrieval, in addition to the specific cues the subject manipulates to facilitate retrieval. The general context is only partly under subject control, since much of it is environmentally determined and even the internally generated context (e.g., what the subject is currently "thinking about") may be difficult to alter radically. Thus, changes in probe cues will be under subject control to only a degree, and this fact is one of the factors underlying retrieval failure.

This extremely brief and superficial overview should not be allowed to give the reader an impression that distinguishing the automatic and controlled phases of retrieval is generally a simple matter. Furthermore, the controlled phase can be a most complex and many-faceted process. To give just one example, Anderson and Bower (1973) consider how subjects compare test sentences with stored sentences in long-term memory. In our model, there are two phases to this experiment: (a) a retrieval of the appropriate informational nexus from long-term memory through the use of probe information (including the test question) and (b) a comparison in short-term memory of the test question and the retrieved search set. The Anderson and Bower model, in the present view, is primarily a model of the second of these processes—the comparison within short-term memory of the test sentence and the retrieved information. See Shiffrin's (1970, 1975b, 1976) studies for a more elaborate discussion of long-term retrieval processes.

V. Models of Search and Detection

In this section the theory we have proposed will be compared with some of the previous models that have been proposed for detection and search tasks.

A. *The Sternberg Model: Serial, Exhaustive Search*

The serial, exhaustive search model is presented in Sternberg (1966). Extensions, further evidence, and reviews of the literature are given in Sternberg (1969a, 1969b, 1975). This model is meant to apply to memory search tasks ($F = 1$, M varies) with small memory-set sizes. It has been confirmed impressively often but is limited in scope, as it is meant to apply only to a small range of paradigms. In our present theory, serial, exhaustive search is considered to be one of the controlled search strategies, one that is very commonly adopted.

The factors that lead subjects in our situation to adopt terminating, controlled search remain uncertain at present. Possibly the much larger search loads in our study led to a terminating strategy. One other minor discrepancy in results involves the Sternberg (1966) finding apparently showing fixed sets (CM) and varied sets (VM) to give equivalent set-size functions. In retrospect we can see that Sternberg's subjects in the fixed set procedure were given far too little training, and the fixed sets were varied too often for automatic detection to have developed.

1. *The Relation Between Sternberg's Model and Our Framework*

Reaction time was proposed by Sternberg (1969a, 1969b) to be the sum of motor response time and the times for four additive stages: (a) encoding (affected by stimulus legibility), (b) serial comparison (affected by size of the memory set), (c) binary decision (affected by whether a positive or negative trial has occurred), (d) translation and response organization (affected by relative frequency of positive and negative trials). In Figure 11, in which our more general framework is depicted, these stages proposed by Sternberg may be placed as follows: "Encoding" is equivalent to automatic encoding; "serial comparison and binary decision" both are part of controlled processing; "response organization and motor response execution" are part of response production. In conclusion, then, the Sternberg model differs from the detailed quantitative model we fit to our data in Part I, but nevertheless may be viewed as a special case of our general framework.

2. *Problems for the Serial, Exhaustive Comparison Model and Their Resolution*

The serial, exhaustive model is particularly important on account of the research it has generated that give results not consistent with the theory in its simplest form. These inconsistencies have led to many new models, some of which will be discussed below. The apparent inconsistencies include the following results.

1. Reaction times depend on the serial presentation position of the tested item within the memory set (in VM procedures with memory sets changing every trial). The most pronounced effects occur when the memory-set items are presented quickly and the display item is presented very soon after the last memory-set items (Burrows & Okada, 1971; Clifton & Birenbaum, 1970; Corballis, 1967; Klatzky, Juola, & Atkinson, 1971; Klatzky & Smith, 1972).
2. Reaction times depend on the relative frequency of presentation of items within the memory set in CM procedures (see Biederman & Stacy; Krueger, 1970; Miller & Pachella, 1973; Shiffrin & Schneider, 1974; Theios et al., 1973).
3. Reaction times can be speeded for cued items in VM paradigms (Klatzky & Smith, 1972) or speeded for expected items in CM paradigms (Shiffrin & Schneider, 1974; but note that the fixed sets changed every 160 trials so that only low degrees of practice were involved), or speeded in a VM paradigm for an item repeated during the pre-

sentation of the memory set (Baddeley & Ecob, 1973).

4. The outcomes in CM paradigms often differ in fundamental ways from those in VM paradigms, especially when some simple physical basis separates the memory and distractor sets, or when the degree of training is high. Many such findings have been discussed extensively in earlier sections of this paper and will not be reviewed again here.

5. Subjects typically give incorrect responses on 1–10% of the trials, depending upon the condition. The serial search models do not generally posit explicit mechanisms to predict errors.

Before turning to alternative models, it is useful to consider how these phenomena are dealt with in the framework of our theory, and also to see how Sternberg explains the results.

The various findings in CM paradigms are easiest to explain: Automatic detection develops that enables the serial search to be bypassed. Much of the research in this paper was directed toward establishing this fact. Sternberg (1975) is less specific but also suggests that some alternative, more efficient, search process is used in such situations.

Within varied-mapping paradigms, recourse to the hypothesis of automatic detection cannot be used to explain the results. It would be parsimonious if each of the above findings (1) to (3) proved to be the result of a common mechanism. Shiffrin and Schneider (1974) proposed one possible mechanism—namely, that when information concerning test probabilities is available to the subject, then one item is placed in a special state prior to each test display (called a state of "expectancy"), and that a test of the expected item results in a faster response time than tests of nonexpected items. Stimulus probability, serial position, cueing, stimulus repetition, or differential importance could all be expected to determine the probability with which stimuli will be expected.

The effects of expectancy could possibly act at several different stages en route to the execution of the response. One possibility, also mentioned by Sternberg (1975), is that the comparison process is carried out in a speeded fashion for an expected item. Another possibility is that the encoding process or the response production stage is speeded when an expected item is tested. Shiffrin and Schneider (1974) attempted to carry out a test discriminating among the possible models. While the general notion of expectancy was supported by the results, the study was not conclusive in determining which version of the model was to be preferred.

The important point to be emphasized is that we (and Sternberg also) suggest that findings (1) to (3) listed above are not incompatible with a controlled search process that is serial and exhaustive. We suggest these findings can be explained by factors that affect other stages of the response process than the comparison stage or perhaps by a factor that affects the comparison process only on certain trials. Many other researchers have preferred to discard the hypothesis of a serial, exhaustive comparison process. Some of the models proposed as alternatives will be discussed below.

The final problem for the serial, exhaustive search model is posed by the occurrence of errors. The explanations of controlled search processes by Sternberg (and by us) for reaction time tasks do not propose an explicit mechanism by which errors may occur. Many investigators have simply ignored errors as long as the error rate rate is low, under 5%, say. This is perhaps justifiable if errors are anomalous events that do not interact with any of the other variables being studied. Indeed the stability of findings across numerous studies that do not attempt to fix error rates at any given value and that vary considerably in the observed rates (in the range 0–10%) lends some support to the view that ignoring errors will not distort the conclusions drawn from the reaction time data.

In principle, however, error rates cannot be ignored. Pachella (1974) has shown that instructions used to change error rates by just a few percent can cause considerable changes in the level and pattern of reaction times. There are many ways in which error-producing mechanisms can be appended to controlled search models. Two of the most important are as follows: First, the subject might terminate, or be forced to terminate, the controlled search before the search is completed. Then a response would have to be made as a guess based on partial information at most, or perhaps no response at all would be made (an omission). These types of error resulting from early search termination account for almost all errors in search and attention tasks using above-threshold stimuli with accuracy as a measure; in particular, most of the errors in the accuracy tasks reported in the present paper and in Part I are errors of the type. Second, the search process might be carried out incorrectly; that is, a comparison might be carried out incorrectly owing to confusions among items, forgetting, misperception, and the like. When error mechanisms are appended to the basic search model a new expanded theory results, which must be tested through joint consideration of error data and reaction time data.

Two basic approaches may be utilized to test search theories incorporating error predictions. One approach involves manipulating the error rate systematically within each condition; the manipulation may be carried out through instruction, deadline training, or signals-to-respond (Reed, 1973). This approach has the advantage of making full use of both the error and reaction time data from a single experiment. It has the disadvantage that the manipulations of error rates may affect the nature of the controlled search strategy adopted by the subject. For example, Reed (1976) carried out a remarkable series of tests of a wide variety of models using data collected in the signal-to-respond procedure. Unfortunately, though one model could be singled out in preference to the others, the basic data differed in important respects from those ordinarily found in the simpler version of the same paradigm.

The alternative approach involves carrying out two separate studies involving the same subjects. One study utilizes the usual reaction time methodology, with instructions to keep error rates low. The second study utilizes a paradigm in which the available search time is systematically varied and the error rates corresponding to each amount of available search time are collected. The model derived from one study can then be used to predict the results of the other. This is the method that was adopted in Experiments 1 and 2 of Part I. This approach has the advantage of relating two fields of inquiry that are normally treated separately. In the present paper, for example, attention tasks using accuracy as a measure were linked to search tasks using reaction time as a measure.

It should be noted, by the way, that the type of error predicted for the results of Part I/Experiment 1 is that caused by premature termination of controlled search whose characteristics were derived from the reaction time results of Part I/Experiment 2. However, the errors seen in Part I/Experiment 2 were not necessarily caused by the same mechanism. In fact many, if not most, of them may have resulted from confusions, forgetting, or anomalous condition-independent factors.

B. *The Theios Model*

There are really two separate treatments to be discussed here: the specific micromodel used by Theios et al. (1973) to predict results that would otherwise be fit by a serial, exhaustive model, and the general theory used by Theios (1973, 1975) to describe the production of response times in a variety of tasks.

1. *Serial, Terminating Search Through a List of Memory Items and Distractors*

In this model a list is constructed in memory, a list on which appear all items that might be displayed for test. Each of these items has an associated response cue (positive or negative) attached to it. The memory-set items and distractors may in general be intermingled in this list, but a variety of factors affect list ordering, and in many cases the memory-set items may tend to appear early in the list. During the test, the subject serially scans the list until the test item is located, then terminates the search and responds according to the cue information located along with that item. A group of items at the end of the list is assumed, however, to be searched in parallel, in one step, if search has not previously terminated. These last items are supposed to be accessed through LTS, while those scanned serially are in STS. (There are a variety of complications and extensions added to this model that we shall not discuss.) In at least some situations, this model can approximately predict linear set-size functions that are parallel for negative and positive trials.

Various versions of the model were fitted to data collected by Theios et al. (1973). They utilized a CM procedure with a moderate amount of training. It seems likely, therefore, that many of the observed effects in their study, especially those based on differences in stimulus frequency, resulted from the development of at least a small degree of automatic detection. The more frequent items would come to elicit automatic responses sooner, causing the observed reduction in response reaction times. In our terms, the subject's strategy was probably a mixture of automatic detection and controlled search, making simple conclusions difficult. Nevertheless, even if automatic detection is not involved, it may be asked whether models of the complexity suggested by Theios (or suggested by the comments above) are necessary to fit his data. Shiffrin and Schneider (1974) showed that a simple version of the expectancy model using serial, exhaustive search through the positive set could provide a fit to the data about as good as that provided by the more complex models of Theios et al. (1973). Under these circumstances the data do not provide strong support for the Theios model.

Aside from questions concerning the details of the Theios et al. (1973) study, one can address the more general hypothesis put forward by these investigators. This hypothesis suggests that most memory search studies are best conceived of in terms of a (rather complex) serial, terminating search model. Although we have collected data arguing for a serial, terminating search through the memory set (Experiment 2, Part I), we agree with Sternberg (1975) that a serial, terminating search is not a good candidate as a model for the simpler paradigms in which linear, parallel set-size functions are collected (see Sternberg, 1975, for the relevant arguments). There are many additional reasons for us to reject the hypothesis of Theios et al; these reasons arise from the results of the present paper. Basically, the Theios et al. model was proposed to deal with a variety of stimulus probability effects that arise in CM paradigms. We have shown in this article that indeed a different detection process develops in such paradigms; however, the fully developed automatic process and the pure controlled process are two different mechanisms and both are relatively simple. In effect, Theios et al. have been led to propose a rather complex unitary model to try to fit data that probably reflect a mixture of two different simple search mechanisms.

2. *A Hybrid Serial-Parallel Model*

A much more general theory of response time generation in search tasks has been presented by Theios (1973, 1975). This theory has many similarities to the theory presented in the present article. It postulates input and identification stages (which we would call automatic encoding), a response determination stage (which in our terms is either a controlled search or an automatic process or some combination, depending on the paradigm), and response program selection and response output stages (which we term response production). Theios's descriptions of the situations in which automatic processing should occur are quite similar to those we suggest in this paper. His distinction between controlled and automatic processing is not made as explicit as in the present article, and in his view the use of controlled search is tied more to stimulus–response compatibility than to the type and amount of practice (i.e., the consistency of the mapping). Furthermore, the details he presents concerning controlled search processes are somewhat different and more limited than we have presented in the present treatments. On the whole, however, the Theios (1975) theory is quite compatible with the general approach taken in this paper; in fact, we suggest that interested readers read that paper for additional evidence regarding automatic processing in reaction time tasks.

C. *Parallel Processing Models*

Certain parallel processing models have been proposed for VM search paradigms as alternatives to the usual serial models. Atkinson et al. (1969), Murdock (1971), and Townsend (1971) have all presented versions of such models. The Townsend model can serve as an example. In this model the time to complete any one comparison is distributed exponentially with a rate constant that depends on the number of current items undergoing processing. When any item completes comparison, the rates are immediately readjusted.

Such models almost completely mimic serial models. The same set-size predictions are derived as for serial models, whether or not termination of search is assumed. There are some minor problems with these models. For example, the exponential assumption implies a particular relation between the growth of the mean reaction time and the growth of the variance of the reaction time, with set size. More important, however, this sort of parallel model is not conceptually very different from the serial models—comparisons occur, on the average, at evenly spaced intervals of time for each of the items of the memory set (or the display set). We prefer, therefore, to think of such models as members of a class also containing the serial models. As Sternberg (1966) has shown, some types of parallel models, in which each item is processed independently of the number of alternative items, cannot predict the observed results.

D. *Direct Memory Access Models Based on Trace Strength*

A number of models have included a search mechanism that resembles in certain respects the automatic detection process presented in this paper. In these models, an item presented for test is encoded directly to some location in long-term memory, a location containing information enabling the correct response to be given. These models are most applicable to CM paradigms, although they have sometimes been applied to VM situations and although some have been elaborated to contain serial, short-term search as an additional process. Generally, in these models the trace strength of the code in long-term memory determines the ease of access and the response time for a given item.

Speaking generally, it is our feeling that such models capture part of the process we have described as automatic detection. Note, however, that the two concepts are not identical. An example will make this point clear. Suppose we define the memory set to consist of all words whose fourth letter is alphabetically prior to the second letter. A word presented auditorally will presumably be encoded automatically until the long-term memory node is located that contains the spelling. Then a controlled process will check the spelling to see whether the criterion is satisfied. Such a process is not "automatic detection." If several items are presented at once, each would have to be checked serially. This is an example of a task designed so that the information found in long-term store corresponding to a given input is sufficient for a decision to be made correctly, without considering other members of the memory set. This property could hold for either VM or CM tasks, regardless of task-specific training and regardless of whether an attention (or other task-specific) response has been learned.

The automatic detection mode is therefore similar to the search mode of trace-strength models in that each input is encoded automatically, and relevant information is located in LTS. If the information located must be processed in limited, controlled fashion (so that multiple inputs would have to be decided about sequentially, for example) then we would not consider the process to be one of automatic detection. On the other hand, if the information located is attention directing, or even response eliciting, then the process would be one of automatic detection. Several varieties of trace-strength models may be distinguished. We consider these next.

1. *Trace-Strength Discrimination Models*

In these models, the information found in memory for any tested item consists of a unidimensional value of strength (often interpreted as familiarity). The task is designed so that the response can be based on the retrieved strength value, and it is assumed that criteria can be chosen so that errors are kept low in frequency. Examples of such models are found in Corballis, Kirby, and Miller (1972); Nickerson (1972); Baddeley and Ecob (1973); Cavanagh (1973, 1976), and Anderson (1973).

There are many tasks in which we regard a response strategy based on trace strength to be highly plausible, whether the trace strength is utilized in a controlled decision or used to initiate an automatic response. But we must ask whether such models are reasonable when applied to typical VM search paradigms. There is no question that the models are capable of generating accurate predictions if strengths for items in variously sized memory sets are appropriately adjusted. However, a model with this much freedom could predict anything and would not be interesting. In practice, of course, various assumptions are made to govern the possible changes in trace strengths across conditions. The extant models, however, do not appear to us to have the simplicity and elegance of serial search theory when both are applied to VM search tasks.

The models of Corballis (1975) and Cavanagh (1976) both assume that trace strength is affected by rehearsal (called "sequential priming" by Corballis). Both models have been developed primarily in response to the findings of serial position effects. As discussed earlier, such effects are not incompatible with serial search, or even serial, exhaustive search. Thus if rehearsal affects search in important ways it might do so because it controls the trace strengths used to respond (in the trace-strength models), or because it affects serial search order (in a terminating search), or because it affects encoding or response time rather than comparison time (in a serial, exhaustive model). Future research will be necessary to establish which, if any, of these possibilities is true.

2. *The Atkinson-Juola Model*

This model (Atkinson & Juola, 1974) was constructed to explain the results of search studies in which a large, previously learned ensemble of items forms the memory set. The model assumes that the test item is evaluated in terms of its familiarity—a value above a criterion leads to a positive response, a value below a lower criterion leads to a negative response, and a value between the two criteria leads the subject to carry out a serial search of the list. Such a model is quite compatible with our theory as applied to a task of this kind, but there is at least one problem requiring further research. The estimated serial search rate is only about 10 msec per item. If the serial search mode were similar to that in the usual procedure then a rate nearer 40 msec would be expected. There are several hypotheses to explain the discrepancy—for example, it is possible that the search is somehow restricted to a subset of items that are similar to the test item. A somewhat different hypothesis would suggest that no serial search occurs at all, but that some rechecking of the familiarity value is necessary when the first observation is between the criteria, and that the confusability of the test item, and hence the need to recheck the test item, will depend on the list length.

In paradigms like that of Atkinson and Juola (1974), where trace strength or familiarity is used as a basis for the response (thereby sometimes bypassing a serial search), it might be asked to what degree the detection is an automatic process. Certainly the encoding of the stimulus and the generation of a familiarity value is an automatic process. The decision how to respond for a given value of familiarity could very well be a controlled process initially, though with sufficient practice in the task, the decision and response initiation for extreme familiarity values might also become automatized.

E. *Some Conclusions about Search Models*

Each of the models we have considered has some elements in common with our general theory. At the same time, each has been applied to some paradigms that we think would be better handled by an alternative process. The direct access trace-strength theories, for example, have a close affinity with our automatic detection mechanism and have a natural application to CM paradigms, tasks with perceptual cues separating the memory and distractor sets, or tasks in which associations in LTS to each test item contain information enabling a response to be made. Such models should not, however, be applied, in our view, to VM paradigms of the type that are usually studied in the laboratory. For such tasks serial, controlled search seems a more attractive possibility (or at least limited controlled search of one sort or another). Much of the present article has been devoted to demonstrating that two qualitatively different detection or search modes exist and to establishing the conditions under which each is utilized. Many of the models to date have attempted, we think unsuccessfully, to apply a single mode of search to all the search paradigms.

The second point we wish to emphasize is that many of the prior models have been constructed to deal with just a few studies or just a few results from those studies. The history of investigation in this area has demonstrated beyond doubt that numerous models are capable of predicting results from any given study. A proper evaluation of models should incorporate two tests: (a) Can the model predict a wide variety of results in differing paradigms, and can it predict results in both the reaction time and accuracy domain? (b) Can the model predict results from a single series of studies on the same subjects, a series in which most of the commonly examined variables are manipulated?

VI. Models of Selective Attention

A. *The Broadbent and Treisman Models*

The Broadbent (1971) model supposes that there are two basic selective processes, one leading from the physical input to a set of internal codes that serve as the evidence, and a second leading from the internal codes to categories and responses. The selection in Stage 1 is called "filtering" or "stimulus set," and that in State *2* is called "pigeonholing" or "response set." There is a very short-term sensory-information store (less than 1-sec duration) called a "buffer," which accepts information in parallel from the physical inputs. (This store is equivalent to the more peripheral information in STS in our theory.) The filter selects information from the buffer and sends it through a limited-capacity channel. In the Broadbent (1958) theory the filter allowed information from only one source at a time to continue through the system. In the theory proposed in 1971, Broadbent's earlier view was modified in accord with data put forth and theory suggested by Treisman (e.g., 1960, 1969). This modified theory suggests that the filter does not block, but only attenuates, information from nonattended sources. The states of evidence in the system are the output of the filter.

While there are obvious elements of similarity between our approach and Broadbent's, we would like to focus on the differences:

1. *The Role of Control Processes*

In our theory, selective attention is largely defined in terms of control processes. Selectivity due to structural characteristics of the processing system, even structure that has been learned, is not considered to be an attentional process. In the Broadbent model, the distinction between the structural, learned components of selectivity and those under subject control are not as clearly distinguished. One consequence of Broadbent's approach is that the type and amount of training is not implicated as a particularly important determinant of the presence or absence of attentional deficits. Much of the data in the present article, however, has shown how the conditions of practice (particularly the consistency of the mapping of attention to a feature to a reinforced outcome) determine the development of automaticity, the bypass of controlled search, and the elimination of attentional limitations.

The emphasis on the role of control processes in our theory also implies that we are advocating an active rather passive view of attentional selectivity. In our theory attentional selectivity is largely the result of accentuation of certain informational elements through the use of control processing; in the Broadbent-Treisman approach, on the other hand, inhibition and attenuation of the processing of certain inputs play the most important role in selectivity.

2. *Selectivity Leading to the Internal States of Evidence*

In our theory the internal states of evidence are the result of automatic processing of sensory inputs, with subject control affecting sensory encoding only indirectly and to a small degree. On the other hand, the selectivity that results from filtering plays a large and fundamental role in Broadbent's theory. This fact alone would not be a source of discrepancy between the theories if these filtering processes reflected only relatively permanent, structural components of selectivity. However, the examples of filtering given by Broadbent make it clear that this process is intended to reflect temporary shifts in attentional control. Thus, in numerous studies, filtering is said to select inputs on the basis of a physical difference, like location. In some of these studies the same physical cue is always assigned consistently, so that automatic detection (in our terms) could have developed, but the number of trials is usually so small that controlled processing probably was utilized. In other of these types of studies the physical cue was assigned randomly across trials, so that selection must have been based on temporary control processes (see Broadbent, 1971, chap. 5). In either class of studies, Broadbent suggests filtering as the basis for selection, with attended inputs presumably resulting in a better state of internal evidence than nonattended inputs.

We take a different view. We propose that automatic processing results in roughly equivalent internal states, but that controlled processing must examine these internal states sequentially. Thus information on an attended channel is examined first, and is therefore detected and recalled more effectively than information on a nonattended channel. We will discuss below some studies by Shiffrin and his colleagues that seem to provide strong evidence in favor of this viewpoint.

3. *Short-Term Store and Levels of Processing*

In the Broadbent theory there is a short-term buffer more or less corresponding to a short-term memory for peripheral information. In the original 1958 theory this was the sole basis for short-term memory, but in the 1971 theory, a later short-term system called "primary memory" was added subsequent to the filter and the limited-capacity channel. Thus the new theory is in close correspondence with the approach of Atkinson and Shiffrin (1968).

However, the Broadbent theory identifies selectivity with particular *levels* of processing: Selectivity operates after the buffer and prior to primary memory. We, on the other hand, treat all short-term memory as a single continuum consisting of the results of automatic encoding as well as inputs from LTS. Thus STS consists of information at a wide variety of levels of processing. In our view, selectivity is not restricted to any special levels. Rather, selectivity operates at whatever level controlled processing is utilizing. If controlled processing is checking color serially, then attentional selectivity will appear at the peripheral level of hue, while if words are compared serially to see which is a synonym of a memory item, then attentional selectivity will appear at the central semantic level.

4. *The Role of Timing*

An important concept in Broadbent's theory is that of "switching time." Broadbent interprets the results of many studies by assuming that attention (the filter) may not be redirected from one source to another before the lapse of some minimal time, called switching time. Although some of the results Broadbent ascribes to switching time limitations are due in our theory to other mechanisms, we have no objection to arguments that attentional limitations are rooted in limitations on the rate of processing operations. To the contrary, we have carried this argument much further and have suggested that the time utilized during controlled processing is in many tasks the immediate, proximal cause of the inability to divide attention. The results of Experiments 1 and 2 in Part I, and our analysis of them, were intended to demonstrate the role of timing by relating the accuracy of performance in attention tasks to the reaction times produced in search tasks.

B. *The Shiffrin (1975a) Model*

The work by Shiffrin and his colleagues summarized by Shiffrin (1975a) may be regarded as a preface to the present studies and present theory. The primary aim of the earlier studies was a demonstration that encoding quality is virtually unaffected by attentional instructions. The results from a series of experiments demonstrated that information from a source (location) is processed as well when it is the only information requiring processing as when it arrives simultaneously with other information also requiring processing. If filtering or attenuation were occurring between stimulus input and the production of internal states of evidence, then attention would have been allocated less effectively when simultaneous inputs were used, and performance would have suffered.

It could well be argued that the results were obtained only because attentional capacity had not been exceeded. In effect, this is our own argument. Such an argument does not help resuscitate the views that early filtering, attenuation, or blocking of processing takes place, since the situations in which the results were obtained were as complex at the stimulus input side as most of the studies in which attentional deficits are found.

It should be noted that the Shiffrin (1975a) results are easily accommodated in Broadbent's general theory. It must be assumed that the various inputs are retained in the sensory buffer long enough that they may all be passed successfully through the limited-capacity system (in essence this is our explanation, also, if "sensory buffer" is replaced by "STS"). The features in the sensory buffer must be high-level ones, however, because in the studies masks were used to delete low-level features from storage. Thus such results call into question the need for two selective processes, one before and one after generation of states of internal evidence. We argue instead that only one controlled selection process occurs, after automatic encoding has produced states of internal evidence. Even if it should eventually be determined that some small degree of selection occurs during initial perceptual processing, the accumulation of data makes it clear that the magnitude of attentional effects due to controlled processing after automatic encoding is far greater than the magnitude of any effects due to early selection. For further discussion see Shiffrin & Gardner (1972), Shiffrin, Craig, & Cohen (1973), Shiffrin, Gardner, & Allmeyer (1973), Shiffrin & Grantham (1974), Shiffrin, Pisoni, & Casteneda-Mendez (1974), and Shiffrin, McKay, & Shaffer (1976). The hypothesis supported by these studies, that of a single selection phase following automatic encoding, has been proposed by a number of researchers and theorists. We turn now to these proposals.

C. *The Deutsch and Deutsch Theory*

As preface to the Deutsch and Deutsch (1963) model and to other similar theories, we will review some results in auditory selective perception from shadowing or dichotic listening tasks. Then we will consider the Deutsch and Deutsch treatment, objections to it, and our present view.

1. *Experiments on Auditory Shadowing and Dichotic Listening*

Starting with Cherry's (1953) study, numerous studies have utilized a shadowing technique in which the subject repeats continually a stream of speech (usually presented to one ear) while other messages are presented simultaneously (usually to the other ear). Moray (1959) showed that information on the nonshadowed ear could not be recalled, recognized, or relearned with savings. However, the subject's own name on the nonshadowed ear could be noticed and recalled at the end of the experiment. The traditional explanations for such results hold that the information presented to the nonattended ear is greatly attenuated, but that some information, enough to let highly salient inputs be noticed, does pass the filter. The basis for the types of information that are noticed on the nonshadowed ear is usually considered to be rooted in the content of the input. The earlier treatments supposed that simple physical cues served as a basis (not only in Broadbent, 1958, but also in Neisser, 1967). For example, two messages in the same voice are confused, but if the messages are in a man's voice and a woman's voice, then the nonshadowed message can be ignored. A number of studies by Treisman (see 1969 for a review) showed that content of messages can also serve as a basis for selection though usually not as effectively as physical cues.

2. *The Deutsch and Deutsch Approach*

Faced with selection on the basis of content, Deutsch and Deutsch (1963) suggested that all inputs are analyzed to a relatively high level and that the results of the processing are then used to select certain stimuli for further processing, for memory, and for response. We are in essential agreement with this view, though we have elaborated on it considerably and introduced the important distinction between controlled search and learned attention responses leading to automatic detection.

The Deutsch and Deutsch model has not yet gained general acceptance and in particular has been rejected by Broadbent (1971) and Treisman and Riley (1969). Within the framework of shadowing studies, several objections have been raised to the Deutsch and Deutsch position. First, selection without a physical cue remains difficult even with wide variations in content. Second, differences in content affect selection very little when a physical cue is present. Third, selection of a given word in a dichotic sequence depends on the previously attended word and not on the previously presented, but unattended, word.

Our present theory, with its distinction between controlled search and automatic detection, overcomes these objections. If the stimuli presented do not cause automatic-attention responses, then a controlled search situation (like that of Experiment 4a) obtains, and a similar explanation holds: The subject will carry on attention-demanding, controlled processing primarily on the shadowed or attended message, with only minimal controlled processing of the nonattended information, just enough to establish which information is to be given deeper processing. Physical cues, when available, will make this relevancy decision easy. On the other hand, if automatic-attention responses have been preexperimentally attached to stimuli, these stimuli will be attended to and remembered even when they are presented on a to-be-ignored channel. This fact explains the recall for the subject's own name, for example. Furthermore, the objection that selection is difficult without a physical cue does not hold, since we have shown that consistent training leads to the development of automatic-attention responses even without the presence of a simple, consistent physical cue (though a great deal of training may be necessary—see Experiments 1 and 3).

To return to the original controversy, some of the most persuasive evidence in favor of the Deutsch and Deutsch position has been the subject's ability to divide attention without deficit, even in VM situations when automatic detection could not have been operating. Such data have often been discounted by Broadbent and others on the following basis. It is argued that relatively unanalyzed information is held temporarily in a sensory buffer; thus it should not be surprising that multiple stimuli can pass through the filter as long as a free period exists after presentation that allows the information to be passed serially from the buffer through the filter. This argument is used to explain the apparent ability to analyze multiple simultaneous stimuli in a variety of tasks and is thereby used to discount much of the evidence supporting the Deutsch and Deutsch position. Broadbent's argument, however, does not account for the results of our multiple-frame studies, for in these studies there were no blank periods to allow serial processing, and yet the rate of presentation could even be speeded under the consistent (CM) training conditions. We do not disagree that under VM conditions a postpresentation free period will allow simultaneous inputs to be handled serially by the attention system. We merely wish to point out that this explanation cannot help explain the relationship between our VM and CM conditions in the multiple-frame tasks. Instead, Broadbent and Treisman would have to argue that the salience of the consistently trained targets is so high that the filter plays no role in processing whatsoever, in which case it might as well be argued that automatic processing bypasses the selective system.

In summary, we argue that the position of Deutsch and Deutsch, when extended in the manner of our present theory, is both fully defensible and in better accord with the data than the opposing views.

D. *The Norman Model*

Norman, alone and with coauthors, has been responsible for a number of rather different models. We consider here the version most closely in accord with that of Deutsch and Deutsch (Norman, 1968, 1969; see also Lindsay & Norman, 1972). This model is basically an extension of the Deutsch and Deutsch approach in the sense that it proposes a relatively complete analysis of all incoming stimuli. As in the Deutsch and Deutsch approach, a mechanism is proposed by which certain of these complete encodings are selected for further serial processing. This mechanism is called "pertinence."

In the model, each set of encodings is assigned a value of pertinence that determines the order of further processing. The value of pertinence is assumed to derive from two sources: the stimulus encoding and an analysis of previous inputs. The item selected for further analysis is that one whose pertinence value is highest as determined by a combination of sensory encoding and previous analysis. In very general terms this model is highly similar to the theory we have proposed.

However, the Norman model fails to make the important and necessary distinction between automatic-attention learning and controlled processing. Thus extrapolating from Norman's theory, visual search for a single memory item, even in a VM task, should be almost capacity-unlimited and parallel. The subject should, by hypothesis, raise the pertinence value of the single item, and hence this item should be processed first. However, all of our studies have demonstrated that such pertinence changes cannot take place on a trial-by-trial basis. It is only after considerable CM training that an automatic-attention response develops (or, in Norman's terms, that pertinence is raised). Thus a subject may *wish* to find a particular stimulus but does not in general have the attentional control to enable the stimulus to be located without a serial, limited, controlled search.

E. *The Neisser Model*

Neisser (1967) has proposed a general model that postulates an early stage of parallel processing not under subject control called "preattentive" (similar to our systemic, automatic stage), and a later, controllable, serial stage referred to by the term "focal attention." In many respects this general approach is quite similar to the present treatment. For example, Neisser discusses in some detail the fact that preattentive processes can control attention or responses (as in eye movements). Furthermore, it is suggested that practice can lead to search becoming dependent on preattentive mechanisms, and hence lead detection to become automatic.

Neisser's own treatment of his visual search data differed in some respects from our present treatment. For example, Neisser argued that irrelevant targets in CM designs are not processed in as much detail as relevant targets. We argue that automatic analysis is equal, but that attention is drawn to the relevant targets. Furthermore, there are some differences in emphasis between Neisser's and our models. For example, Neisser tends to emphasize that preattentive discriminations are based on relatively crude physical differences among stimuli, whereas we feel "preattentive" processes are determined by conditions of training and practice, with physical differences affecting the rate of development of automaticity. However, such differences should not be allowed to mask the similarity of the general approaches.[9]

F. *The LaBerge Model*

LaBerge (1973, 1975) presents an attentional model, in the context of reaction time matching tasks, that has many elements in common with the present theory. He supposes, like Deutsch and Deutsch, that there exists an automatic, learned encoding system that operates in parallel on input stimuli and produces features independently of controlled, attentive processing. Depending on prior learning, there are no limits to the depth of such processing. Subsequent to such automatic encoding, controlled processing requiring attention is necessary to carry out further analyses of the input. The features resulting from automatic encoding are assumed to be capable of calling attention to themselves. LaBerge proposes, and demonstrates empirically, that new automatic encodings, that is, perceptual learning, may occur with practice. Finally, LaBerge emphasizes that perceptual learning takes place *mainly after high accuracy is achieved.* All of these aspects of the system are quite similar to ours.

One additional element of the LaBerge model holds that attention may be directed toward a not-yet-learned feature, causing it to act temporarily as an automatically learned unit, but only while attention is maintained. That is, a feature that would normally require extended controlled processing to be encoded from automatically activated, simpler subunits, can, when attention is focused on it in advance, be activated by the perceptual encoding process. This suggests at least some degree of control over encoding. While we have downplayed the role of such control, we have not ruled it out. In fact, we suggested that the results from Part I/Experiment 2, when M was equal to 1, might well be explained by such a mechanism.

If there is a difference between our approach and LeBerge's, it lies in the different areas of applications and the completeness of the assumptions. For example, we have elaborated on the conditions of training that determine the ability of features to attract attention. That is, we have shown how the attention-attracting ability of features depends on the use of consistent-mapping training procedures.

On the whole, our model and LaBerge's agree when applied to the same data. Thus, although the two theories have been built on somewhat different data bases, they are consistent even to the degree that there are similarities in the diagrammatic representations (compare our Figures 11 and 12 with the depictions of the model in LaBerge, 1973, 1975).

Let us now return to models that are in basic agreement with the Broadbent-Treisman view, as opposed to the Deutsch and Deutsch view.

G. *The Kahneman Theory*

Kahneman (1973) proposes a rather extensive theory and reviews the literature in considerable detail. We cannot begin to do justice to his overall system in the limited space available and will therefore discuss only a few general features. Kahneman proposes a general theory of the allocation of capacity according to a factor called "effort." In our terms, he is concerned with the general limits on controlled processing: How can control be allocated? Can total capacity vary? What types of operations can be carried out together with what losses of efficiency?

Interpreting our theory in Kahneman's terms, one would say that we have proposed a specific time-based theory determining controlled processing capacity in detection tasks, and that our theory gives a set of rules or

---

[9] One must take care to distinguish the Neisser (1967) view of attention from the Neisser (1967) view of cognitive processing in general. The general view emphasizes the constructivist aspects of perception, sometimes known as "top-down" processing, in which perception is heavily influenced at all stages by the subject's expectations and general knowledge. The essence of this general view seems to be a bit incompatible with the notion of an initial automatic processing system that is fixed, permanent, and largely unaffected by higher level controlled processing.

strategies by which subjects reallocate their processing in such tasks. Such a time-based theory is not specifically given by Kahneman. Furthermore, in Kahneman's theory there are many results taken to indicate a change in allocatable capacity, a change often determined by changes in effort. We feel that some of these findings are better explained by recourse to the hypothesis that automatic processing has developed. As we have argued in this paper, the development of automatic processing provides a means of improving performance independently of the effort put into controlled processing. In fact, the effort required is usually greatly reduced by the development of automatic processing even though performance improves.

There is one other important specific area of difference between Kahneman's theory and ours. In common with a number of the theories we have described (i.e., Broadbent), Kahneman's assumes that control of allocation affects processing prior to the stage at which features indicating recognition are activated. We have dealt with this issue at length and will not repeat the arguments, but we feel that control does not operate at this point in the midst of what we call automatic processing (except in certain special cases, such as instances where the inputs are degraded or ambiguous, in which case automatic encoding never reaches the recognition stage).

Reversing the emphasis, it may be asked what elements in Kahneman's approach are not considered in our theory. The principal such element is the entire concept of effort. We have confined ourselves to tasks in which we assume effort is continually maintained at the highest possible level. Of course it must be true that performance will drop when effort drops to a low enough level, but our theory has not addressed this question at all. It is natural, however, for us to assume that changes in effort will affect controlled processing but not automatic processing.

H. *The Norman and Rumelhart Model*

The attentional features of this model (Norman & Rumelhart, 1970) derive largely from Rumelhart's (1970) article. This model is aimed to describe visual detection in threshold tasks. It differs from many previous models in that it is quantitative and hence makes explicit, testable predictions. It makes a key assumption that the number of features abstracted, and the rate of features abstracted, from a given input stimulus will be inversely related to the total number of items simultaneously presented. That is, attention must be shared (i.e., divided) among multiple inputs, even at the earliest stages of processing. The simultaneous–successive studies by Shiffrin and Gardner (1972), described earlier, most clearly show that this assumption is false. The features abstracted are of equivalent worth whether or not the inputs are successive or simultaneous. As we have seen in the present article, the difficulties in dividing attention arise later, during controlled processing, subsequent to the encoding of features.

I. *Classification Models for Attentional Paradigms*

A number of investigators have provided useful and interesting classifications of attentional paradigms. One example is a classification of types of attentional limitations by Norman and Bobrow (1975). In general, Norman and Bobrow suppose that inputs may be given limited processing for either of two reasons: "data-limitations" and "resource-limitations." In our terms, data-limitations refer to the performance-reducing interactions that occur during automatic encoding; for example, a signal is heard better in quiet than in a background of white noise. Resource limitations refer to the kind of attention limitations that we would classify as the result of controlled processing. Norman and Bobrow argue that many authors have given insufficient thought to the source of the performance limitations observed in their tasks. They argue that many tasks can be varied parametrically (quantitatively) so that data-limitations occur under some conditions and resource limitations under other conditions. We agree fully with these arguments; our studies have, we think, spanned the range of these limitations and demonstrated these points quite clearly.

Many other investigators have provided classifications of attentional paradigms and results. The Treisman (1969) review, for example, is largely atheoretical and classifies a variety of paradigms and tasks with a view toward indicating the types of attention restriction that occur in each. Similarly Moray (1969a, 1969b) has classified attentional studies and paradigms. Unfortunately, the variety of results and the experimenter's ability to design new tasks seem to have grown about as fast as the growth in complexity of the classifications. Thus as useful as these classifications have been, they have been limited to a posteriori descriptions.

J. *Other Models*

The brief summary of a number of theories given above by no means exhausts the field. Keele (1973) and Hochberg (1970) presented models like that of Deutsch and Deutsch and generally similar to ours. Important and influential reviews, with considerable theoretical import, have been written by Egeth (1967), Moray (1969a, 1969b), Lindsay (1970), and Swets and Kristofferson (1970). Important research has been carried out by Eriksen and his associates (e.g., Eriksen & Collins, 1969; Eriksen & Eriksen, 1974; Eriksen & Hoffman, 1972; Eriksen & Rohrbaugh, 1970) and by Posner and his associates (e.g., Posner & Boies, 1971; Posner & Snyder, 1975). Although space limitations preclude our reviewing this work, we feel that the models we have discussed above represent a cross section of typical approaches that have been put forward for tasks that resemble those in this paper.

We would like to finish this discussion of alternative approaches by considering very briefly the relation of our theory to those arising in the fields of animal discrimination learning and attention (e.g., see Mackintosh, 1975; Sutherland & Mackintosh, 1971). We feel that data and theory in the human and infrahuman areas are in much closer accord than previous work might lead one to believe. There has been a tendency to compare attentional models deriving from controlled processing effects in varied-mapping studies in humans to models deriving from automatic-attention learning in consistent-mapping paradigms in animals. The results in the present article should make it clear that such a correspondence will fail. In fact, however, it is becoming clear that a comparison of VM with VM, and CM with CM, studies in the two areas reveals similar findings. To give just one example, Riley and Leith (1976) review some paradigms in which animals are forced to utilize what we term controlled processing; in such cases they reveal capacity limitations and attentional control roughly similar to that found in human subjects (see also Mackintosh, 1975, and Wagner, 1976, for discussions of related matters).

K. *Summary*

Many of the elements constituting our theory may be found in previous proposals. There are certain themes, which have recurred in earlier theories, about which we take a particular position.

1. *Control of processing during encoding and feature abstraction.* We argue that control is minimal, that subject-controlled filtering, blocking, or attenuating does not occur during this stage but subsequent to perceptual encoding.
2. *The role of crude physical differences in leading to automatic discrimination.* We argue that the conditions of training are crucial to the development of automatic processing and automatic detection; the nature of physical differences among items is of secondary importance theoretically, though it may in practice be important in determining the rate of acquisition and levels of performance during acquisition.
3. *The limitations on, and capacity of, controlled processing.* We have presented a specific, detailed model in which limitations on rate of processing determine the controlling system's capacity.

On the one hand, our theory may be seen as both a specification and extension of views proposed in various earlier models. On the other hand, our theory is grounded in a much firmer empirical foundation (i.e., Experiments 1–3 of Part I, and 1–4 of Part II) than the previous proposals, encompasses a greater breadth than some of the earlier views, and states its assumptions precisely enough that at least in our standard paradigm and in a number of others as well, predictions can be made in advance of experimental manipulations. Perhaps most important, we have, in Part I, taken steps toward quantitatively linking attentional phenomena and search mechanisms, and have thereby tied together two important fields of inquiry—attention and search—that have too often been treated separately in previous models.

VII. Summary and Conclusions

The studies reported in this article were designed to test and extend the ideas presented in Part I. In Part I we showed that controlled search is utilized in varied-mapping search paradigms and that automatic detection is utilized in consistent-mapping search paradigms. These findings obtained in both multiple-frame tasks measuring accuracy and in single-frame tasks measuring reaction time. The VM results in both types of tasks were fitted by a quantitative model assuming serial, terminating search, with comparisons cycling through the frame for each memory-set item.

The experiments in the present article were designed to examine the learning and unlearning of automatic detection, the role of categorization, and the learning of automatic attending. A summary of the main findings are given below.

Experiment 1 was carried out to study the development of automatic processing. Automatic detection for Letter Set 1 in a background of Letter Set 2 developed over several thousand training trials during which time performance improved considerably. Then the two sets were reversed and performance not only deteriorated sharply but dropped below the level obtaining at the start of training when controlled search was being utilized. The results showed that

1. Automatic detection could develop when the memory and distractor sets were not preexperimentally categorized.

2. The learning of an automatic-attention response is a long-term phenomenon greatly resistant to change (since many more trials were needed to recover to a given performance level after reversal than to reach that level in initial training).

3. The presence among the distractors of items that automatically attract attention harms the subject's ability to carry out a controlled search (since postreversal performance dropped below initial performance levels).

4. Eventually automatic-attention responses can be "unlearned" and new sets of automatic responses learned, but only after considerable amounts of retraining.

Experiment 2 was carried out to test the effects of reversal when the two sets involved were well-known categories. In Experiment 2, the subjects who had been trained extensively to search for digits in a background of consonants (or vice versa) were reversed in a fashion similar to that used in Experiment 1. The results showed

1. Performance in a VM condition after reversal was identical to that in normal VM controlled search, even though all items presented were from the previous CM target category. Hence, when all items have equally strong automatic-attention responses, these either "cancel each other out" or can be ignored, and normal controlled search takes place.

2. A reversal of memory and distractor sets for the CM conditions reduced performance almost to chance levels. This result shows that the phenomenon of automatic detection is not due to the presence of a categorical difference between memory and distractor sets, since a distinction between letters and numbers was still available after reversal.

3. The pattern of results after reversal indicated that a controlled search was being utilized, but one that compared the category of each input to the memory-set category in a single operation. The results suggest that categories do improve search performance but do so by facilitating controlled search.

Experiment 3 was designed to test directly the role of categories in the development of automatic detection and the operation of controlled search. It was shown that in VM conditions, category learning for arbitrary, visually confusable sets of letters eventually took place after 25 sessions of training. At this point, controlled search was still being used, but the comparison process had switched from individual items to categories as a whole (memory-set size had no effect). However, when training was then switched to a CM procedure, performance improved radically and automatic detection was learned. The results suggest that

1. Both categories and automatic responses can be learned for arbitrarily constituted, visually confusable character sets.
2. The role of categories is to improve controlled search, and possibly to speed the acquisition of automatic detection.

As a group, Experiments 1, 2, and 3 provided a convincing demonstration of the qualitatively different characteristics of controlled search and automatic detection.

Experiment 4 was designed to test the subjects' ability to focus attention in the face of distraction by (a) neutral characters, (b) current targets in to-be-ignored display positions, (c) items in to-be-ignored positions that subjects had been trained previously to respond to with automatic-attention responses. The results showed that controlled search can usually be directed to locations that the subject desires to attend to but that automatic-attention responses can overwhelm the controlled processing system and can cause attention to be allocated to positions that should be ignored. Thus automatic-attention responses cannot be ignored and will interrupt and redirect ongoing controlled processing.

The attention literature as a whole was considered, and in light of our results the following rules were suggested.

1. Divided-attention deficits arise from limitations on controlled processing. In particular, detection deficits are due to the limited rate of the serial comparison process.
2. Dividing attention is possible when the targets have been consistently mapped during training until automatic detection operates.
3. Focused-attention deficits arise when the distracting stimuli initiate automatic-attention responses.
4. Focusing attention is possible during controlled processing (e.g., in VM tasks).

We shall conclude by summarizing briefly the progress we feel has been made in organizing the phenomena of detection, search, attention, and perceptual learning. One of the greatest contributions of the work is the fact that all these areas have been tied both empirically and theoretically to each other and to common mechanisms. Another important contribution is the delineation of the differences between, and the characteristics of, automatic and controlled processing. We make no claims that these experiments have concluded the study of these areas. Quite to the contrary, one of the most important results of this work is the host of new questions that can now be phrased, and we hope, answered.

Reference Notes

1. Sperling, G. *Multiple detections in a brief visual stimulus: The sharing and switching of attention.* Paper presented at the 16th annual meeting of the Psychonomic Society, Denver, November 1975.
2. Sperling, G. *Attention operating characteristics for visual search.* Paper presented at the 9th annual Mathematical Psychology meeting, New York, August 1976.
3. Shiffrin, R. M., & Cook, J. R. *Short-term forgetting without rehearsal or interference: Data and theory.* Paper presented at the 16th annual meeting of the Psychonomic Society, Denver, November 1975.

References

Anderson, J. A. A theory for the recognition of items from short memorized lists. *Psychological Review,* 1973, *80,* 417–438.

Anderson, J. R., & Bower, G. H. *Human associative memory.* Washington, D.C.: V. H. Winston & Sons, 1973.

Atkinson, R. C., Holmbren, J. E., & Juola, J. F. Processing time as influenced by the number of elements in a visual display. *Perception & Psychophysics,* 1969, *6,* 321–326.

Atkinson, R. C., & Juola, J. F. Search and decision processes in recognition memory. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 1). San Francisco: Freeman, 1974.

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.

Atkinson, R. C., & Shiffrin, R. M. The control of short-term memory. *Scientific American,* 1971, *225,* 82–90.

Baddeley, A. D., & Ecob, J. R. Reaction time and short-term memory: Implications of repetition effects for the high-speed exhaustive scan hypothesis. *Quarterly Journal of Experimental Psychology,* 1973, *25,* 229–240.

Biederman, I., & Stacey, E. W., Jr. Stimulus probability and stimulus set size in memory scanning. *Journal of Experimental Psychology,* 1974, *102,* 1100–1107.

Bjork, R. A. Short-term storage: The ordered output of a central processor. In F. Restle et al., (Eds.), *Cognitive theory* (Vol. 1). Hillsdale, N.J.: Erlbaum Associates, 1975.

Brand, J. Classification without identification in visual search. *Quarterly Journal of Experimental Psychology,* 1971, *23,* 178–186.

Bransford. J. D., & Franks, J. J. The abstraction of linguistic ideas. *Cognitive Psychology,* 1971, *2,* 331–350.

Broadbent, D. E. *Perception and communication.* London: Pergamon Press, 1958.

Broadbent, D. E. *Decision and stress.* London: Academic Press, 1971.

Burrows, D. & Okada, R. Serial position effects in high-speed memory search. *Perception & Psychophysics,* 1971, *10,* 305–308.

Cavanagh, P. Holographic processes realizable in the neural realm: Prediction of short-term memory performance (Doctoral dissertation, Carnegie-Mellon University, 1972). *Dissertation Abstracts International,* 1973, *33,* 3280B. (University Microfilms No. 72–234, 79)

Cavanagh, P. Holographic and trace strength models of rehearsal effects in the item recognition task. *Memory & Cognition,* 1976, *4,* 186–199.

Cherry, C. Some experiments on the reception of speech with one and with two ears. *Journal of the Acoustical Society of America,* 1953, *25,* 975–979.

Cherry, E. C., & Taylor, W. K. Some further experiments upon the recognition of speech with one and two ears. *Journal of the Acoustical Society of America,* 1954, *26,* 554–559.

Clifton, C., & Birenbaum, S. Effects of serial position and delay of probe in a memory scan task. *Journal of Experimental Psychology,* 1970, *86,* 69–76.

Clifton, C., & Gutschera, K. D. Evidence for hierarchical search processes in a recognition memory task. *Journal of Verbal Learning and Verbal Behavior,* 1971, *10,* 528–541.

Corballis, M. C. Serial order in recognition and recall. *Journal of Experimental Psychology,* 1967, *74,* 99–105.

Corballis, M. C. Access to memory: An analysis of recognition times. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V.* New York: Academic Press, 1975.

Corballis, M. C., Kirby, J., & Miller, A. Access to elements of a memorized list. *Journal of Experimental Psychology,* 1972, *94,* 185–190.

Corteen, R. S., & Wood, B. Autonomic responses to shock-associated words in an unattended channel. *Journal of Experimental Psychology,* 1972, *94,* 308–313.

Craik, F. I. M., & Jacoby, L. L. A process view of short-term retention. In F. Restle et al., (Eds.), *Cognitive theory* (Vol. 1). Hillsdale, N.J.: Erlbaum Associates, 1975.

Craik, F. I. M., & Lockhart, R. S. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior,* 1972, *11,* 671–684.

Craik, F. I. M.. & Tulving, E. Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General,* 1975, *104,* 268–294.

DeRosa, D. V., & Morin, R. E. Recognition reaction time for digits in consecutive and nonconsecutive memorized sets. *Journal of Experimental Psychology,* 1970, *83,* 472–479.

Deutsch, J. A., & Deutsch, D. Attention: Some theoretical considerations. *Psychological Review,* 1963, *70,* 80–90.

Egeth, H. Selective attention. *Psychological Bulletin,* 1967, *67,* 41–57.

Eriksen, C. W., & Collins, J. F. Temporal courses of selective attention. *Journal of Experimental Psychology,* 1969, *80,* 254–261.

Eriksen, B. A., & Eriksen, C. W. Effects of noise letters upon identification of target in nonsearch task. *Perception & Psychophysics,* 1974, *16,* 143–149.

Eriksen, C. W., & Hoffman, J. E. Some characteristics of selective attention in visual perception determined by vocal reaction time. *Perception & Psychophysics,* 1972, *11,* 169–171.

Eriksen, C. W., & Rohrbaugh, J. W. Some factors determining efficiency of selective attention. *American Journal of Psychology,* 1970, *83,* 330–342.

Estes, W. K. Interactions of signal and background variables in visual processing. *Perception & Psychophysics,* 1972, *12,* 278–286.

Gibson, E. J. *Principles of perceptual learning and development.* New York: Prentice-Hall, 1969.

Gleitman, H., & Jonides, J. The cost of categorization in visual search: Incomplete processing of targets and field items. *Perception & Psychophysics,* 1976, *20,* 281–288.

Gordon, G. T. Interaction between items in visual search. *Journal of Experimental Psychology,* 1968, *76,* 348–355.

Gould, J. D., & Carn, R. Visual search, complex backgrounds, mental counters, and eye movements. *Perception & Psychophysics,* 1973, *14,* 125–132.

Green, B. F., & Anderson, L. K. Color coding in a visual search task. *Journal of Experimental Psychology,* 1956, *51,* 19–24.

Hochberg, J. E. Attention, organization, and consciousness. In D. I. Mostofsky (Ed.), *Attention: Contemporary theory and analysis.* New York: Appleton-Century-Crofts, 1970.

Homa, D. Organization and long-term memory search. *Memory & Cognition,* 1973, *1,* 369–379.

Hyde, T. S., & Jenkins, J. J. Differential effects of incidental tasks on the organization of recall of a list of highly associated words. *Journal of Experimental Psychology,* 1969, *82,* 472–481.

Hyde, T. S., & Jenkins, J. J. Recall for words as a function of semantic, graphic, and syntactic orienting tasks. *Journal of Verbal Learning and Verbal Behavior,* 1973, *12,* 471–480.

Ingling, N. W. Categorization: A mechanism for rapid information processing. *Journal of Experimental Psychology,* 1972, *94,* 239–243.

Jensen, A. R., & Rohwer, W. D. The Stroop Color-Word Test: A review. *Acta Psychologica,* 1966, *25,* 36–93.

Jonides, J., & Gleitman, H. A conceptual category effect in visual search: O as letter or as digit. *Perception & Psychophysics,* 1972, *12,* 457–460.

Jonides, J., & Gleitman, H. The benefit of categorization in visual search: Target location without identification. *Perception & Psychophysics,* 1976, *20,* 289–298.

Jung, J. Comments on Mandler's "From association to structure." *Psychological Review,* 1965, *72,* 318–322.

Kahneman, D. *Attention and effort.* Englewood Cliffs, N.J.: Prentice-Hall, 1973.

Kahneman, D. Effort, recognition, and recall in auditory attention. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V.* New York: Academic Press, 1975.

Keele, S. W. Attention demands of memory retrieval. *Journal of Experimental Psychology,* 1972, *93,* 245–248.

Keele, S. W. *Attention and human performance.* Pacific Palisades, Calif.: Goodyear, 1973.

Klatzky, R. L., Juola, J. F., & Atkinson, R. C. Test stimulus representation and experimental context effects in memory scanning. *Journal of Experimental Psychology,* 1971, *87,* 281–288.

Klatzky, R. L., & Smith, E. E. Stimulus expectancy and retrieval from short-term memory. *Journal of Experimental Psychology,* 1972, *94,* 101–107.

Krueger, L. E. Effect of stimulus probability on two-choice reaction time. *Journal of Experimental Psychology,* 1970, *84,* 377–379.

LaBerge, D. Attention and the measurement of perceptual learning. *Memory & Cognition,* 1973, 1, 268–276.

LaBerge, D. Acquisition of automatic processing in perceptual and associative learning. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V.* New York: Academic Press, 1975.

Lewis, J. L. Semantic processing of unattended messages using dichotic listening. *Journal of Experimental Psychology,* 1970, *85,* 225–228.

Lindsay, P. H. Multichannel processing in perception. In D. I. Mostofsky (Ed.), *Attention: Contemporary theory and analysis.* New York: Appleton-Century-Crofts, 1970.

Lindsay, P. H., & Norman, D. A. *Human information processing: An introduction to psychology.* New York: Academic Press, 1972.

Lively, B. L., & Sanford, B. J. The use of category information in a memory search task. *Journal of Experimental Psychology,* 1972, *93,* 379–385.

Mackintosh, N. J. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review,* 1975, *82,* 276–298.

Mandler, G. From association to structure. *Psychological Review,* 1962, *69,* 415–426.

Mandler, G. Subjects do think: A reply to Jung's comments. *Psychological Review,* 1965, *72,* 323–326.

Miller, J. O., & Pachella, R. G. Locus of the stimulus probability effect. *Journal of Experimental Psychology,* 1973, *101,* 227–231.

Moray, N. Attention in dichotic listening. Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology,* 1959, *11,* 56–60.

Moray, N. *Attention: Selective processes in vision and hearing.* New York: Academic Press, 1969. (a)

Moray, N. *Listening and attention.* Harmondsworth, Middlesex, England: Penguin Books, 1969. (b)

Moray, N. A data base for theories of selective listening. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V.* New York: Academic Press, 1975.

Murdock, B. B., Jr. A parallel-processing model for scanning. *Perception & Psychophysics,* 1971, *10,* 289–291.

Naus, M. J. Memory search of categorized lists: A consideration of self-terminating search strategies. *Journal of Experimental Psychology,* 1974, *102,* 992–1000.

Naus, M. J., Glucksberg, S., & Ornstein, P. A. Taxonomic word categories and memory search. *Cognitive Psychology,* 1972, *3,* 643–654.

Neisser, U. Decision time without reaction time: Experiments in visual scanning. *American Journal of Psychology,* 1963, *76,* 376–385.

Neisser, U. *Cognitive psychology.* New York: Appleton-Century-Crofts, 1967.

Nickerson, R. S. Binary-classification reaction time: A review of some studies of human information processing capabilities. *Psychonomic Monograph Supplements,* 1972, *4,* 275–318.

Norman, D. A. Towards a theory of memory and attention. *Psychological Review,* 1968, *75,* 522–536.

Norman, D. A. Memory while shadowing. *Quarterly*

Journal of Experimental Psychology, 1969, *21,* 85–93.

Norman, D. A. & Bobrow, D. G. On data-limited and resource-limited processes. *Cognitive Psychology,* 1975, *7,* 44–64.

Norman, D. A., & Rumelhart, D. E. A system for perception and memory. In D. A. Norman (Ed.), *Models of human memory.* New York: Academic Press, 1970.

Okada, R., & Burrows, D. Organizational factors in high-speed scanning. *Journal of Experimental Psychology,* 1973, *101,* 77–81.

Pachella, R. G. The interpretation of reaction time in information processing research. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition.* Hillsdale, N.J.: Erlbaum Associates, 1974.

Posner, M. I., & Boies, S. J. Components of attention. *Psychological Review,* 1971, *78,* 391–408.

Posner, M. I., & Snyder, C. R. R. Facilitation and inhibition in the processing of signals. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V.* New York: Academic Press, 1975.

Postman, L. Short-term memory and incidental learning. In A. W. Melton (Ed.), *Categories of human learning.* New York: Academic Press, 1964.

Postman, L., & Senders, V. Incidental learning and generality of set. *Journal of Experimental Psychology,* 1946, *36,* 153–165.

Reed, A. V. Speed-accuracy tradeoff in recognition memory. *Science,* 1973, *181,* 574–576.

Reed, A. V. List length and the time course of recognition in immediate memory. *Memory & Cognition,* 1976, *4,* 16–30.

Riley, D. A., & Leith, C. R. Multidimensional psychophysics and selective attention in animals. *Psychological Bulletin,* 1976, *83,* 138–160.

Rumelhart, D. E. A multicomponent theory of the perception of briefly exposed visual displays. *Journal of Mathematical Psychology,* 1970, *7,* 191–218.

Schneider, F. W., & Kintz, B. L. An analysis of the incidental-intentional learning dichotomy. *Journal of Experimental Psychology,* 1967, *73,* 85–90.

Schneider, W., & Shiffrin, R. M. Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review,* 1977, *84,* 1–66.

Shiffrin, R. M. Memory search. In D. A. Norman (Ed.), *Models of human memory.* New York: Academic Press, 1970.

Shiffrin, R. M. Information persistence in short-term memory. *Journal of Experimental Psychology,* 1973, *100,* 39–49.

Shiffrin, R. M. The locus and role of attention in memory systems. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V.* New York: Academic Press, 1975. (a)

Shiffrin, R. M. Short-term store: The basis for a memory search. In F. Restle et al. (Eds.), *Cognitive theory* (Vol. 1). Hillsdale, N.J.: Erlbaum Associates, 1975. (b)

Shiffrin, R. M. Capacity limitations in information processing, attention, and memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Memory processes* (Vol. 4). Hillsdale, N.J.: Erlbaum Associates, 1976.

Shiffrin, R. M., Craig, J. C., & Cohen, U. On the degree of attention and capacity limitations in tactile processing. *Perception & Psychophysics,* 1973, *13,* 328–336.

Shiffrin, R. M., & Gardner, G. T. Visual processing capacity and attentional control. *Journal of Experimental Psychology,* 1972, *93,* 72–82.

Shiffrin, R. M., Gardner, G. T., & Allmeyer, D. H. On the the degreee of attention and capacity limitations in visual processing. *Perception & Psychophysics,* 1973, *14,* 231–236.

Shiffrin, R. M., & Geisler, W. S. Visual recognition in a theory of information processing. In R. L. Solso (Ed.), *Contemporary issues in cognitive psychology: The Loyola Symposium.* Washington, D.C.: V. H. Winston & Sons, 1973.

Shiffrin, R. M., & Grantham, D. W. Can attention be allocated to sensory modalities? *Perception & Psychophysics,* 1974, *15,* 460–474.

Shiffrin, R. M., McKay, D. P., & Shaffer, W. O. Attending to forty-nine spatial positions at once. *Journal of Experimental Psychology: Human Perception and Performance,* 1976, *2,* 14–22.

Shiffrin, R. M., Pisoni, D. B., & Castaneda-Mendez, K. Is attention shared beteen the ears? *Cognitive Psychology,* 1974, *6,* 190–215.

Shiffrin, R. M., & Schneider, W. An expectancy model for memory search. *Memory & Cognition,* 1974, *2,* 616–628.

Smith, S. L. Color coding and visual search. *Journal of Experimental Psychology,* 1962, *64,* 434–440.

Sperling, G. The information available in brief visual presentations. *Psychological Monographs,* 1960, *74*(Whole No. 498).

Sternberg, S. High-speed scanning in human memory. *Science,* 1966, *153,* 652–654.

Sternberg, S. The discovery of processing stages: Extensions of Donder's method. In W. G. Koster (Ed.), *Attention and performance II.* Amsterdam: North-Holland, 1969. (a)

Sternberg, S. Memory scanning: Mental processes revealed by reaction time experiments. *American Scientist,* 1969, *57,* 421–457. (b)

Sternberg, S. Memory scanning: New findings and current controversies. *Quarterly Journal of Experimental Psychology,* 1975, *27,* 1–32.

Stroop, J. R. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology,* 1935, *18,* 643–662.

Sutherland, N. S., & Mackintosh, N. J. *Mechanisms of animal discrimination learning.* New York: Academic Press, 1971.

Swets, J. A., & Kristofferson, A. B. Attention. *Annual Review of Psychology,* 1970, *21,* 339–366.

Theios, J. Reaction time measurements in the study of memory processes: Theory and data. In G. H.

Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 7). New York: Academic Press, 1973.

Theios, J. The components of response latency in simple human information tasks. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V*. New York: Academic Press, 1975.

Theios, J. et al. Memory scanning as a serial self-terminating process. *Journal of Experimental Psychology,* 1973, *97*, 323–336.

Townsend, J. T. A note on the identifiability of parallel and serial processes. *Perception & Psychophysics,* 1971, *10*, 161–163.

Treisman, A. M. Contextual cues in selective listening. *Quarterly Journal of Experimental Psychology,* 1960, *12,* 242–248.

Treisman, A. M. Effect of irrelevant material on the efficiency of selective listening. *American Journal of Psychology,* 1964, *77*, 533–546. (a)

Treisman, A. M. Monitoring and storage of irrelevant messages in selective attention. *Journal of Verbal Learning and Verbal Behavior,* 1964, *3*, 449–459. (b)

Treisman, A. M. Verbal cues, language and meanrelevant messages in selective attention. *Journal of Psychology,* 1964, *77*, 206–219. (c)

Treisman, A. M. Strategies and models of selective attention. *Psychological Review,* 1969, *76*, 282–299.

Treisman, A. M., & Riley, J. G. A. Is selective attention selective perception or selective response? A further test. *Journal of Experimental Psychology,* 1969, *79*, 27–34.

Von Wright, J. M. On selection in visual immediate memory. In A. F. Sanders (Ed.), *Attention and performance* (Vol. 3). Amsterdam: North-Holland, 1970.

Von Wright, J. M., Anderson, K., & Stenman, U. Generalization of conditioned GSRs in dichotic listening. In P. M. A. Rabbit & S. Dornic (Eds.), *Attention and performance V*. New York: Academic Press, 1975.

Wagner, A. R. Priming in STM: An information processing mechanism for self-generated or retrieval-generated depression in performance. In T. J. Tighe & R. N. Leaton (Eds.), *Habituation: Perspectives from child development, animal behavior, and neurophysiology.* Hillsdale, N.J.: Erlbaum Associates, 1976.

Williams, J. D. Memory ensemble selection in human information processing. *Journal of Experimental Psychology,* 1971, *88*, 231–238.

Received August 30, 1976 ■