Journal of Experimental Psychology:
Human Perception and Performance
1993, Vol. 19, No. 4, 845-866

0096-1523/93/$3.00

Information Transfer in Iconic Memory Experiments

Karl R. Gegenfurtner and George Sperling

To report letters from briefly exposed letter arrays, subjects must transfer information from a rapidly decaying trace (iconic memory) to more durable storage. In a partial-report paradigm, we systematically varied the proportion (P) of trials with a long cue delay relative to a short cue delay. Practiced subjects used the same transfer strategy independent of P. Data from a partial-report-plus-masking experiment were used to construct a computational model that accurately predicted partial- and whole-report performance with and without masks. Assumptions: Prior to a cue, subjects attend primarily to the middle row of a three-row display, resulting in nonselective transfer. After the cue, they attend only to the cued row. Transfer rate is the product of iconic legibility (which depends on time and retinal location) and attention allocation (which shifts after a cue). Cumulative transfer is limited by the capacity of durable storage.

When subjects are asked to report all the letters they can see in a brief flash of a letter array, they usually can report only four or five letters. The number of reported letters is independent of the number of displayed letters (when more than about five letters are displayed; e.g., Sperling, 1960). One might therefore infer that the limit on the number of letters reported is due to a limited memory capacity, traditionally called the "span of apprehension" (Külpe, 1904; Wundt, 1899). However, a partial-report procedure demonstrates that subjects are able to store a dozen or more items in a very short-term memory (Sperling, 1960).

In a typical partial-report experiment, a 3 × 3 letter matrix is followed by a cue (e.g., a high-, middle-, or low-pitched tone) that indicates the row of the matrix that the subject has to report. Figure 1 shows results from a partial-report experiment. When the cue occurs at the same time as the letters or shortly afterwards, the subject can report all the letters in the cued row. Because the subject does not know in advance which row will be cued, perfect performance implies that all the items are stored and still available at the time of the cue. When the cue is delayed, partial-report performance decreases, until it finally reaches the level of whole report at cue delays of about 500–800 ms.

The decay of partial-report accuracy with cue delay has been taken as evidence for a second kind of memory, which Neisser (1967) called "iconic memory." Neisser assumed that initially all items are held in iconic memory and, at the time of the cue, the cued letters are transferred into a longer lasting storage.

Durable Storage

The partial-report experiment itself does not prove the existence of two different memories. The cue delay effect might be caused by one type of memory, which decays to four or five letters, and the subject has control over which letters survive. However, experiments with a poststimulus mask (Averbach & Sperling, 1960; Sperling, 1960, 1963) show that there is more than one memory. When a poststimulus mask comes soon after stimulus offset, there is a marked decrease in performance relative to the no-mask control condition. Therefore the storage that is probed in the partial-report experiment is destroyed by the mask. When a poststimulus mask comes later, say, 1 s after the stimulus, it does not influence partial reports at all. The interaction of mask and cue delay implies that there are at least two types of memory. One, iconic storage, has a large capacity, decays rapidly, and is destroyed by a mask following the stimulus. The other storage can hold only a limited number of items but is not affected by masking and seems to have a long lifetime. Following Coltheart (1980), we call the second type of memory "durable storage."

There have been several attempts to discriminate among types of durable storage in partial-report and similar types of experiments (Adelson & Jonides, 1980; DiLollo, 1984; Duncan, 1983; Irwin & Yeomans, 1986; Kaufman, 1978; Loftus, Johnson, & Shimamura, 1985; Mewhort, Campbell, Marchetti, & Campbell, 1981; Scarborough, 1972; Sperling, 1967; Townsend, 1973). Similarly, there have been attempts to discriminate between iconic memory as revealed by partial-report experiments, which use informational measures, and other kinds of visual sensory memory that might be revealed by direct sensory judgments, which use synchrony judgments of visual traces with auditory clicks

Karl R. Gegenfurtner, Howard Hughes Medical Institute and Center for Neural Science, New York University; George Sperling, Psychology Department and Center for Neural Science, New York University.

The experimental work was supported by the U.S. Air Force Office of Scientific Research, Life Science Directorate, Visual Information Processing Program; the theoretical work and preparation of this article were supported by Office of Naval Research, Cognitive and Neural Sciences Division.

A preliminary report of Experiments 1 and 2 was made by Sperling and Gegenfurtner (1988).

Correspondence concerning this article should be addressed to George Sperling, who is now at the Department of Cognitive Science, University of California, Irvine, California 92717. Electronic mail may be sent to sperling@uci.edu.

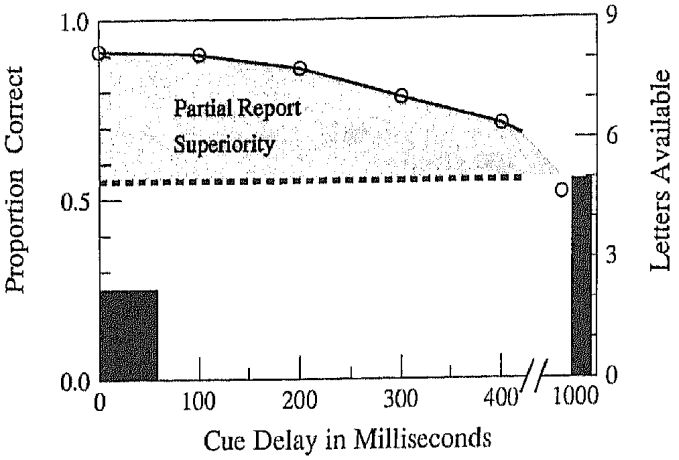

Figure 1. Method and results of a typical partial-report experiment. (The abscissa is the time, relative to the onset of the letter array, of a tonal cue to report a given row. The ordinate is the mean proportion of correctly reported letters in the cued rows. Open circles represent the performance of Subject BL. The filled bar on the right indicates performance in a whole-report experiment. The shaded area between the dashed line and the dots represents partial-report superiority [relative to whole report]. The rightmost data point indicates performance measured at a cue delay of 1,000 ms, which is comparable to the delay involved in the whole-report procedure.)

(Efron, 1970; Sperling, 1967), the integration time between successive stroboscopic flashes (Eriksen & Collins, 1967; Hogden & DiLollo, 1974), and many other procedures (see reviews by Coltheart, 1980; Long, 1980). Because we are concerned only with partial-report experiments here, we can bypass these issues and deal only with iconic memory and durable storage without any further specifications or subdivisions.

Selective and Nonselective Transfer

In a typical iconic memory experiment, at intermediate cue delays, the quality of the iconic image has deteriorated so that few if any additional items can be transferred into durable storage. If the cue were to be further delayed until the iconic image had decayed completely, no transfer from the cued row would be possible, and performance would decrease to 0. As shown in Figure 1, this is not the case. Performance at long cue delays reaches asymptote at whole-report performance, not at zero.

The failure of partial-report accuracy to decay to zero as a function of cue delay is almost certainly due to what Averbach and Coriell (1961) called "nonselective readout." In the time between stimulus and cue, subjects start to transfer items from iconic memory to durable storage. This transfer is nonselective with respect to the cue. It saves subjects from performing badly at long cue delays. On the other hand, the letters transferred nonselectively use up some of the limited capacity of durable storage. By the time the cue appears, durable storage might already be filled with items from the noncued rows. For example, about four items can be transferred in the first 100 ms after stimulus display (Sperling, 1963, 1967). This is approximately the capacity of durable storage. Thus the strategy of nonselective readout at short cue delays might have the disadvantage of overcrowding durable storage, thereby slowing down the subsequent transfer from the cued row.

As several authors (Coltheart, 1980; Dick, 1969; Hall, 1974; Merikle, Lowe, & Coltheart, 1971; Mewhort, Johns, & Coble, 1991; Mewhort, Merikle, & Bryden, 1969; Sakitt, 1976; Sperling, 1960; Sperling & Dosher, 1986) have suggested, it is possible that subjects deliberately use different strategies at different cue delays. On trials with long cue delays, subjects use nonselective readout, to avoid the disaster of having iconic memory decay to illegibility before any items have been transferred. At short cue delays, subjects pay equal attention to all rows and do not transfer items nonselectively, to avoid filling durable storage. Of course, selective strategies would be possible only when partial-report experiments are run in a blocked design, as they typically have been (e.g., Irwin & Yeomans, 1986; Mewhort et al., 1981; Sakitt, 1976; Sperling, 1960). In each block of trials, only one cue delay was used; thereby, subjects can use the strategy that is most advantageous for the particular cue delay.

Overall Plan

In our first experiment, we attempted to discriminate between short-cue-delay and long-cue-delay coding strategies that subjects might use in iconic memory experiments and to determine the costs and benefits of each strategy. It was essential to resolve the possibility that subjects tailor their strategy to the particular cue delay before we proceeded with Experiment 2, a parametric investigation of partial-report accuracy in 25 combinations of Cue Delay × Mask Delay. The data of Experiment 2 enable us to define a model that mathematically describes the time courses of iconic decay and the twin processes of selective and nonselective retrieval. The model, which aggregates all the rows of a three-row stimulus, is very successful computationally, but it contains two enigmas. These enigmas are resolved by noting that characters from the middle row are transferred to durable storage much more rapidly than characters from the other rows and embodying this fact in an elaborated model that treats each row separately. We consider how our model differs from prior computational approaches to iconic transfer. Finally, we show that our formulation of the role of attention in selective transfer is consistent with many other attentional phenomena.

Experiment 1: Coding Strategies

Suppose that subjects in an iconic memory experiment use a strategy of selective transfer at short cue delays and a strategy of nonselective transfer at long cue delays. We wished to estimate the cost of each strategy when it was used inappropriately and the benefit of each strategy when it was

used appropriately in terms of the cost–benefit analysis of Posner, Nissen, and Ogden (1978). The problem in estimating the cost of nonselective transfer at short cue delays was to get the subject to use nonselective transfer (an inappropriate strategy) at short cue delays. Here is the trick. In a blocked situation with only long cue delays, we assumed subjects would certainly use nonselective transfer (the appropriate long-delay strategy). When we included very few trials with a short cue delay in such a block, subjects were still better off from a decision-theoretic viewpoint using nonselective transfer throughout the whole session. Assuming that subjects would try to maximize their performance, we predicted they would use the long-delay strategy in blocks of predominantly long delays and the short-delay strategy in blocks of predominantly short delays. The difference in performance between (a) the few short-delay trials embedded in a block of predominantly long-delay trials and (b) a pure block of short-delay trials provided an estimate of the cost of using nonselective transfer (the inappropriate, long-delay strategy) at short cue delays.

A similar argument was posited for long cue delays. We predicted that when a short cue delay was presented 95% of the time, subjects would use selective transfer (the appropriate strategy). On the occasional long cue delays, we expected their performance to be very poor. Figure 2 illustrates a predicted outcome of this sort of experiment. Performance at long cue delays was expected to decrease markedly with a reduction in the probability of the occurrence of a long cue delay. Performance for both types of cue delay was expected to be highest in the blocked design condition.

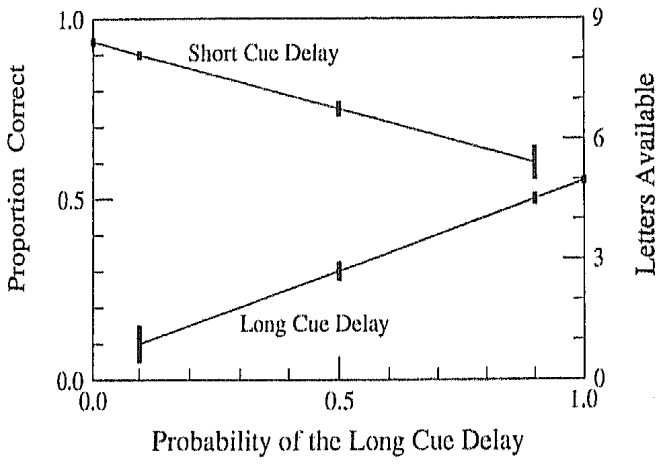

Figure 2. Strategy analysis: The expected outcomes of the cost analysis experiment in which either a long or a short cue delay can occur on a trial. (The abscissa is the probability, within a block of trials, of the long cue delay. The ordinate is the mean proportion of correct reports, conditioned on the type of cue [long vs. short] delay. The upper curve is the expected performance with short cue delays; the lower curve represents long cue delays. Performance at long cue delays is expected to be poor when long cue delays occur rarely [bottom left] and to increase as they become predominant. Performance at short cue delays is symmetrically opposite. The solid bars through the data points indicate standard errors. Their differing length indicates that in the low-probability conditions, fewer trials will be available.)

Method

Subjects. Two graduate and two undergraduate students at New York University participated in the experiment for pay. All subjects had normal or corrected-to-normal vision. Each subject had a minimum of five practice sessions of 200 trials each; for some subjects, practice continued longer until their performance in a regular partial-report experiment with equally likely cue delays reached a steady state. Subjects BL and PC were presented 3 × 3 arrays. Performance for subjects RS and BF was better, so they were shown 3 × 4 arrays.

Stimuli. All experiments were controlled by a Digital Equipment Corporation PDP-11/23 computer. The letters were presented on a Hewlett Packard 1310A cathode ray tube (CRT) with a fast white P4 phosphor. The CRT was driven by a specially designed display interface (Kropfl, 1975) and software for real-time vision experiments (Melchner & Sperling, 1980). Tones were presented on Sennheiser HD414 headphones. A Wavetek Model 159 waveform generator was used to generate the tones, which were set to a comfortable listening level. The timing of the actual stimulus sequences was verified by independent oscilloscopic measurements and was accurate to within 1 ms.

The stimuli consisted of a 3 × 3 or 3 × 4 array of letters. Figure 3a shows a photograph of a typical stimulus. The whole display extended 3.1° or 4.5° of visual angle, respectively, at a viewing distance of 128 cm. Each letter was 1.2 cm high and 1.0 cm wide, with a distance between letters of 2.0 cm horizontally and 1.8 cm vertically. Viewing was binocular.

The luminance of the letters was determined by measuring the luminance of a uniform rectangle with a United Detector Technologies photometer, which had been calibrated against a standard light source. The rectangle had the same pixel intensity as the letters, the same pixel spacing, and the same number of dots as the letter bitmaps. The measured luminance was 34 cd/m^2. The letters were displayed on a dark background of approximately 0.05 cd/m^2. The room was dimly illuminated, and the wall behind the monitor had a luminance of approximately 1.2 cd/m^2. The individual letters were randomly chosen without replacement from the set of 20 consonants, excluding *Y*.

Procedure. Each partial-report session consisted of 200 trials. Figure 4a shows a flow diagram for one trial. The subject initiated the trial by pressing a button. After a random interval of 1.0–1.5 s, the stimuli were displayed for 50 ms (five repeated frames at 10 ms per frame). At the time specified by the cue delay, a tone was sounded on the headphone for 100 ms. The frequencies of the cue tones were 225, 600, and 975 Hz for the bottom, middle, and top row, respectively. The time for the cue delay was measured from the onset of the stimulus. Typically, cue delays in partial-report experiments have been specified in terms of the time from stimulus termination (e.g., Sperling, 1960, and many others). Our reason for specifying a cue delay relative to stimulus onset was that the delay then corresponded to the time for which stimulus information was available before a cue appeared.

Cue delay could be varied independently of the other stimulus parameters, and the cue could occur before stimulus onset, during the stimulus, or after stimulus termination. After the stimulus sequence, the subject was prompted on the screen for a typed response. After the subject responded, the correct letters were shown on the screen, together with the subject's response. Then the next trial started. A response letter was scored as correct only when it was reported in the correct serial position.

In this experimental design, it is inevitable that the low-probability condition for one cue delay coincides with a high probability for the other cue delay. Therefore the number of observations

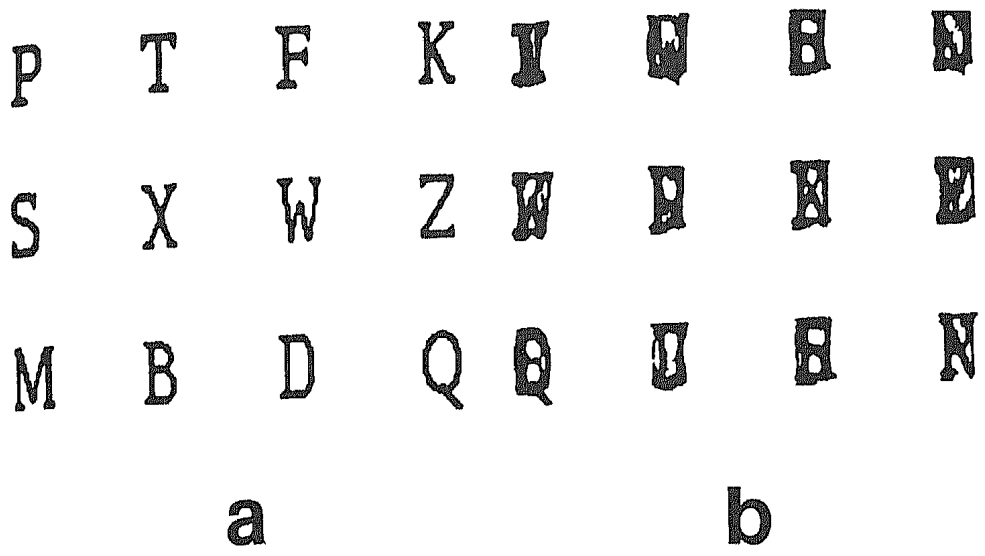

Figure 3. Panel a shows a typical stimulus display. (In both Experiments 1 and 2, all letters were white on a black background; their contrast is reversed here for better reproducibility. Subjects BF and RS used a 3 × 4 matrix of letters; Subjects PC and BL used a 3 × 3 matrix.) Panel b shows a mask. (Like the stimulus, the mask is shown in reversed contrast. Each component masking pattern consists of five different letters displayed in extremely rapid succession.)

for the low-probability conditions is smaller. Performance was evaluated for cue delay probabilities of 0.1, 0.5, 0.9, and 1.0. The short cue delay used was always 0 ms. The duration of the long cue delay was chosen separately for each subject so as to achieve a performance level that would still be better than whole report. The delay values were 400, 800, and 1,000 ms.

Results

The data were analyzed separately for each subject. Figure 5 illustrates the results for subjects PC and BF. Table 1 summarizes the data for all 4 subjects. If there is no effect of probability of occurrence, then performance should not vary and all data points for a fixed cue delay should fall on a straight horizontal line. We therefore estimated slope and intercept of the best fitting lines (in the least squares sense) through the data.

Table 1 shows that the slopes of least-squares-estimated lines through the data points were all negligibly small. Seven of eight slopes were negative, and none of them was significantly different from 0 (according to t tests at the .05 significance level).

Discussion

Three unambiguous aspects of the data lead to three significant conclusions.

Performance in response to a low-probability long-delay cue did not approach zero but reached asymptote at a level typical for whole report. This means that subjects always used nonselective transfer.

Performance in response to a low-probability short-delay cue was not impaired compared with that in response to a high-probability short-delay cue. We infer that nonselective transfer did not involve any additional cost for the subject, even on trials in which selective transfer was also used.

The finding that performance was better for short- than long-delay cues indicates that subjects indeed used selective transfer for short-delay cues.

Finally, with respect to experimental procedures, if subjects have more than one good strategy available, the particular mixture of strategies that they use would depends on the particular mixture of cue delays they confront. The finding that subjects used the same strategy for short- and at long-delay cues greatly simplified the design of Experiment 2. The mixture of cue delays could be optimized for obtaining the desired data, unconstrained by (non)effects on subjects' strategies.

The finding that a single transfer strategy was used at all cue delays is in striking contrast to previous observations suggesting at least two strategies. Sperling (1960, Figure 5) showed a subject whose short-delay-cue strategy failed against long-delay cues and whose long-delay-cue strategy failed to take advantage of short-delay cues. Another (famous) subject (Sperling, 1990, Figure 6) retained his short-delay-cue strategy for too long a cue delay, thereby producing a nonmonotonic iconic decay function. The simplest explanation for the discrepancy between the present data and Sperling's data is that the earlier data were obtained in the first few hundred trials with naive subjects whose performance was clearly nonoptimal. The present data show that, after practice, subjects acquire a single strategy that is effective for both long and short cue delays.

Experiment 2: Time Course of Iconic Memory

The results of Experiment 1 left us with two open questions: How do subjects avoid overfilling durable storage when selective transfer follows nonselective transfer? More generally, how are nonselective transfer and selective transfer combined? To address these questions, we introduced a

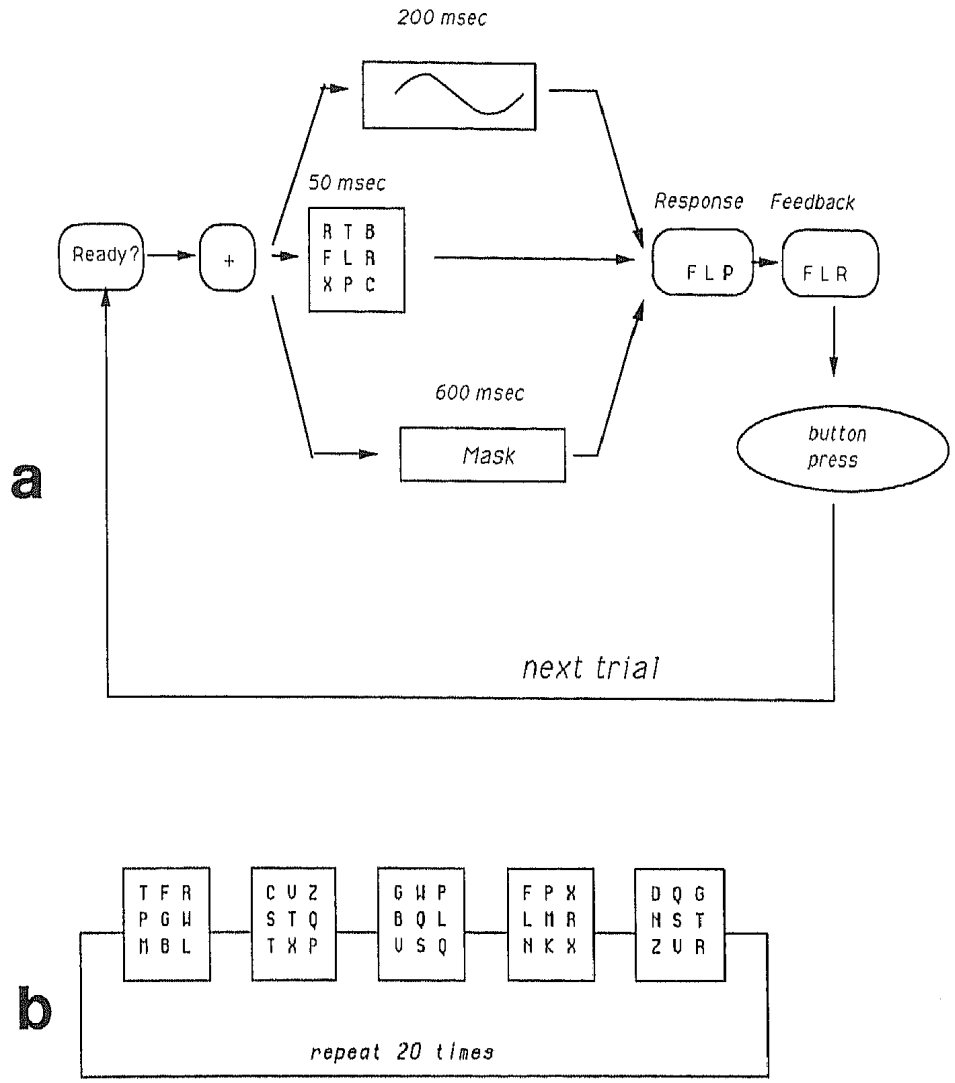

Figure 4. Panel a is a flow chart for a trial. (The three parallel streams for letters, cue, and mask indicate that the onset times for these could be varied independently to produce any arbitrary ordering. The mask was used in Experiment 2 only.) Panel b is a flow chart for the production of a mask. (A sequence of five different frames is painted with 6-ms interframe intervals; the sequence of five is repeated 20 times.)

variably delayed poststimulus mask into the partial-report procedure.

An appropriately chosen visual postexposure masking stimulus should have two properties: It should destroy the contents of iconic memory but leave durable storage unimpaired. For the destruction of iconic memory, a mask is constructed in such a way that when it and the test stimulus are exposed simultaneously, the test stimulus is masked to the point of unintelligibility (Kahneman, 1968; Sperling, 1963). The ability to mask the test stimulus completely when it is strongest (i.e., when it is physically present) implies that the postexposure masking stimulus will even more effectively mask the test after it has been weakened by decay. The ability to leave durable storage unimpaired is demonstrated by showing that long mask delays yield equivalent performance to no-mask control conditions. Such a poststimulus mask serves to limit the time for which information from iconic memory is available for transfer to durable storage. By varying cue delay and mask delay independently in a crossed design, we obtained estimates for the amount of transfer to durable storage in each interval.

Figure 6 illustrates the logic of the masking paradigm in three kinds of conditions. In the first condition (Figure 6a), the cue occurs after stimulus onset and before mask onset. During the interval between stimulus onset and the cue, the subject does not know which row will be cued. Therefore, all transfer is nonselective with respect to the cue. After the cue has occurred, the subject switches attention to the cued row and transfers letters selectively from that row. We call these two kinds of information transfer from iconic memory to durable storage nonselective and selective transfer, respectively.

Two special cases lead to pure selective and pure nonselective transfer. When the mask comes before or at the same time as the cue, only nonselective transfer occurs (Figure 6b). When the cue comes at or before stimulus onset, subjects use selective transfer throughout (Figure 6c). In all other cases (Figure 6a), there is a mixture of selective and nonselective

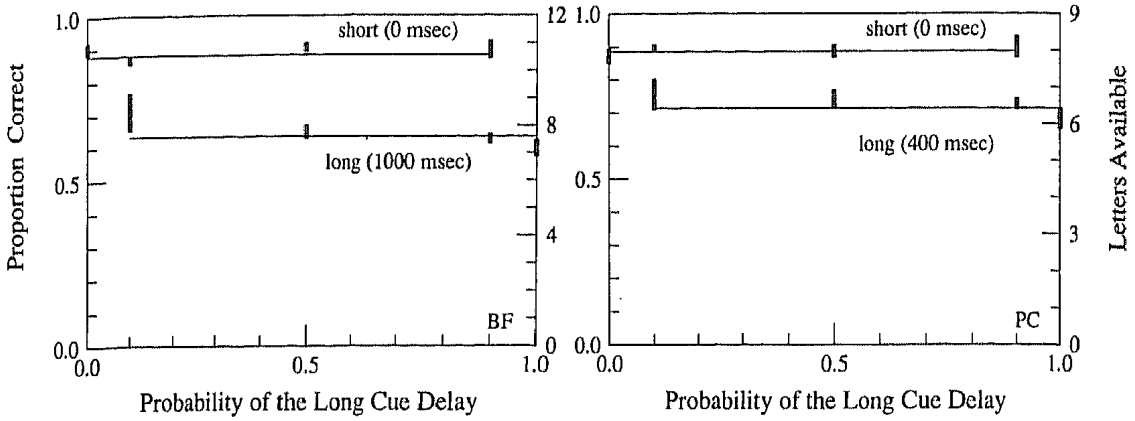

Figure 5. The absence of strategy effects. (Results of Experiment 1 are shown for Subjects PC and BF. The ordinate is the proportion of correct reports; the abscissa is the probability of the long cue delay in a block of 200 trials. The two lines in each panel are the best fitting horizontal lines to the data for each cue delay. The vertical bars represent the standard error [$\pm 1\ \sigma$].)

transfer. Because the cue is irrelevant to nonselective transfer, the pure nonselective conditions should yield the same results as a whole-report experiment with similarly delayed poststimulus masks. This whole-report experiment was carried out as a control condition.

Method

The general experimental methods and subjects were the same as in Experiment 1 except for the following changes.

Subjects. Two subjects, BL and BF, who had served in the Experiment 1, served again in Experiment 2. It should be remarked that BF was able to report one or two items more than average from brief visual exposures. This would place him in the upper 10–20% of subjects in our experience. He was persuaded to serve in this tedious experiment in our hope of discovering some other unusual ability. However, except for a slightly higher level of performance, his data were typical in all respects. In addition, the two other subjects from Experiment 1 served for about half as many trials as BL and BF. Their data did not differ in any important ways from those of Subjects BL and BF and are not presented here.

Masking stimulus. In Experiment 2, a masking pattern (the mask) was displayed at a specified mask delay, which could be shorter or longer than the cue delay. A mask consisted of five different letters displayed in extremely rapid succession at each spatial location, so that the letters were summed by the visual system and could not be recognized individually (Budiansky & Sperling, 1969). All the letters comprising one frame of a mask were painted within 6 ms, a new frame was presented every 6 ms, and the sequence of five different frames was repeated 20 times for a total mask duration of 600 ms. The flow diagram in Figure 4 illustrates this process. The intensity of masks, measured in the same way as the intensity of the letters in Experiment 1, was 47 cd/m^2. Figure 3b illustrates a typical masking pattern. In a brief control experiment, it was verified that recognition of a stimulus letter was at chance when it was presented at the same time as a mask.

Procedure. Mask delays of 100, 200, 300, 400, and 500 ms were used in the experiment. The cue delays chosen were 0, 100, 200, 300, and 400 ms. On each trial, a cue delay and mask delay were chosen randomly in a mixed-list design. Each subject was tested on approximately 5,000 trials in 45-min long sessions of 200 trials each.

Table 1
The Proportion of Correctly Reported Letters as a Function of the Probability of Cue Delays in Experiment 1

Subject/cue delay (ms)	Probability of cue delay 0.1	No. of observations	0.5	No. of observations	0.9	No. of observations	1.0	No. of observations	Slope
BF									
0	0.897	92	0.905	283	0.864	736	0.904	200	−0.015
1,000	0.719	64	0.648	317	0.63	708	0.598	200	−0.012
PC									
0	0.9	40	0.887	140	0.879	360	0.863	200	−0.026
400	0.757	40	0.743	140	0.727	360	0.663	200	−0.082
BL									
0	0.884	23	0.916	103	0.923	181	0.912	80	0.034
800	0.667	19	0.581	97	0.605	177	0.648	53	−0.026
RS									
0	0.983	35	—	—	0.956	244	0.992	32	−0.068
1,000	0.717	23	—	—	0.714	365	0.631	80	−0.06

Note. Slope data indicate the slope of a least squares fitted line through the data points.

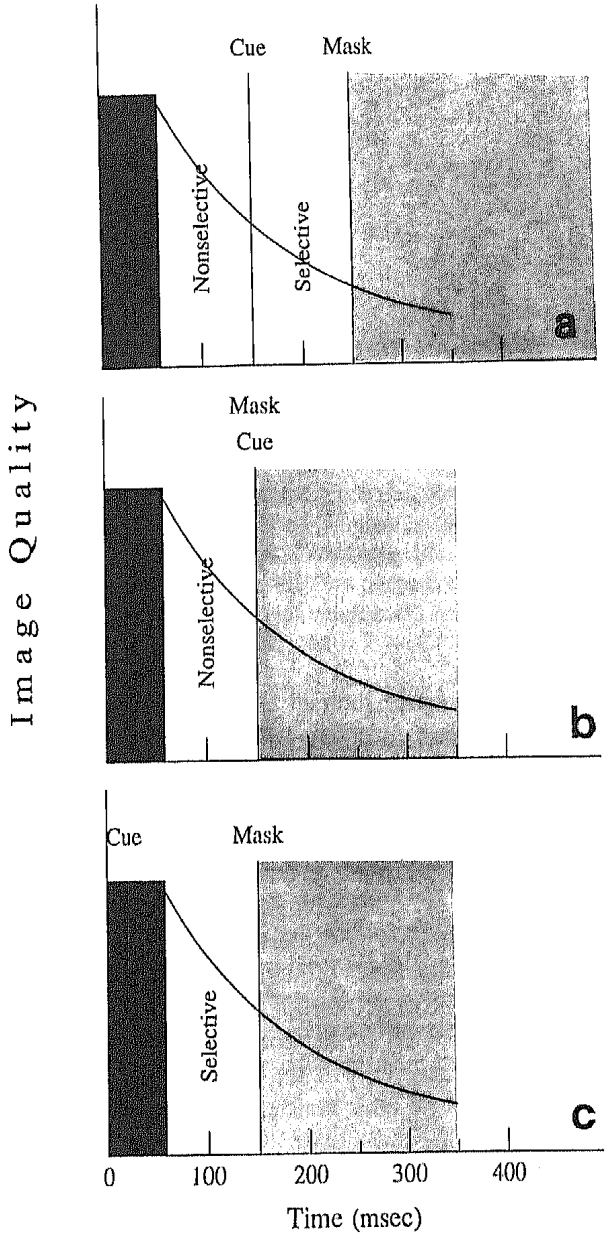

Figure 6. The logic behind the partial-report-plus-mask experiment. (Panel a shows nonselective and selective transfer. The cue occurs before the mask, nonselective transfer occurs before the cue, and selective transfer occurs during the interval cue to the mask. Panel b shows pure nonselective transfer. The cue occurs at or after the onset of the mask. Nonselective transfer ceases after onset of the mask; there is no resumption of transfer after the cue occurs. Panel c shows pure selective transfer: The cue comes at or before the onset of the stimulus.)

Whole report. In the whole-report condition, the subject was asked to report all the letters in the display. The same mask used in the partial-report condition was used. The whole-report practice and test conditions were run in separate sessions after the subjects were already practiced in the partial-report task. Data were collected only after performances had reached asymptote.

Results

As in Experiment 1, the data were analyzed separately for each subject. Figure 7 shows the effect of cue delay with mask delay as a parameter. As in other partial-report experiments, performance dropped as cue delay increased, confirming that subjects made efficient use of the cue. A strictly monotonic decrease in proportion correct as a function of cue delay means that using selective transfer in any time interval yielded more correctly reported letters than did nonselective transfer.

Figure 8 replots the data of Figure 7 with mask delay as the abscissa and cue delay as a parameter. The effect on performance of mask delay was also monotonic; the number of transferred letters increased rapidly with increasing mask delay. A monotonic increase in proportion correct with increasing mask delay means that additional available time for processing the stimulus was always useful.

Pure nonselective transfer. Figure 9 shows data for pure nonselective transfer—all the trials on which the cue occurred simultaneously with or after mask onset. Performance increased very quickly in the first 100 ms and then reached asymptote at around four or five letters. This indicates that these subjects were able to read about four letters in less than 100 ms, which is at the same level that other investigators have found (e.g., Sperling, 1963). Figure 9 also shows the data from the whole-report procedure. Whole-report accuracy is slightly lower than partial-report accuracy. We assume that this slight whole-report deficit was due to the larger number of letters that needed to be reported. Subjects might have occasionally forgotten a letter while reporting the earlier ones. Therefore the partial-report-plus-masking procedure seems to be a slightly better indicator of nonselective transfer than whole report.

The extreme right of Figure 9 shows that whole reports with a 500-ms mask yielded equivalent performance to that in the no-mask control condition. This result indicates that the masking stimulus satisfied the second condition stated for a successful mask: It did not interfere with the contents of durable storage.

Pure selective transfer. The subset of conditions with a cue delay of 0, which indicate pure selective transfer, yielded data that are superficially similar to nonselective-transfer data when graphed in terms of the actual number of letters reported, as shown in Figure 10. Accuracy increased monotonically with mask delay. However, selective transfer took longer than nonselective transfer to approach its asymptotic level (approximately 400 ms vs. 200 ms). The asymptotic accuracy level of selective transfer was much higher than that of nonselective transfer (90% vs. 50%), indicating a partial-report advantage. As in the case of nonselective transfer, when a mask was delayed 500 ms, there was only a negligible difference between mask and no-mask conditions.

An Aggregate-Row Model of Iconic Memory

Experiment 2 characterized purely selective and purely nonselective transfer. In an attempt to explain how they both combine in the overall transfer to durable storage, we developed a model that aggregates performance over rows. Subsequently we found that although the model gave excellent predictions of the present data, it left some serious residual problems. To resolve these, we developed a more

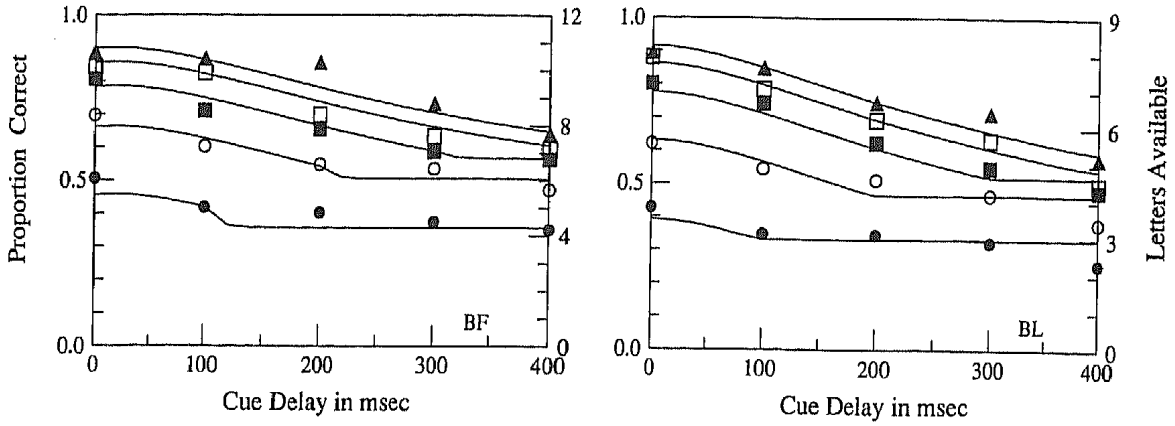

Figure 7. Accuracy of partial reports as a function of cue delay, with mask delay as the parameter: Experiment 2. (The ordinate indicates the proportion of correctly reported letters. The right ordinate indicates the corresponding number of letters transferred to durable storage. Each data point represents 150–250 trials. Panels indicate data for Subjects BL and BF: BL viewed 3 × 3 displays, and BF viewed 3 × 4 displays [Rows × Columns]. The curves drawn through the data points are the best fitting predictions of the two-process aggregate-row model described in the text.)

complicated model in which each row is considered separately. The formulation of the aggregate-row model is presented in this section.

Basic Assumptions: Additivity of Nonselective and Selective Transfer

Both selective and nonselective transfer contribute to the overall performance. In Experiment 2, only the contribution made by nonselective transfer was directly observable. The contribution of selective transfer could be observed only in the absence of nonselective transfer, that is, when selective transfer started immediately at stimulus onset with a cue delay of 0. We now estimate selective transfer at nonzero cue delays. We proceed by making an assumption about the combination rule for selective and nonselective transfer. This assumption allows us to subtract nonselective transfer from overall performance to derive selective transfer at various cue delays.

The simplest combination rule is additivity of the two transfer processes. (Averbach & Coriell, 1961, made a different assumption, which is considered in the Discussion.) To implement additivity of transfer processes, we make the following assumptions. (a) Letters are transferred nonselectively from stimulus onset on until the cue comes. (b) Selective transfer begins at onset of the cue and ends at onset of the mask, when all further information transfer out of iconic memory stops. (c) The total number of letters transferred is the sum of both transfer processes. Specifically, given a cue at time c and a mask at time m, the total number of letters, $L_{c,m}$, transferred from the cued row is the sum of the number of nonselectively transferred letters from the cued row, $(\frac{1}{3})N_{c,m}$, and the number of selectively transferred letters from the cued row, $S_{c,m}$:

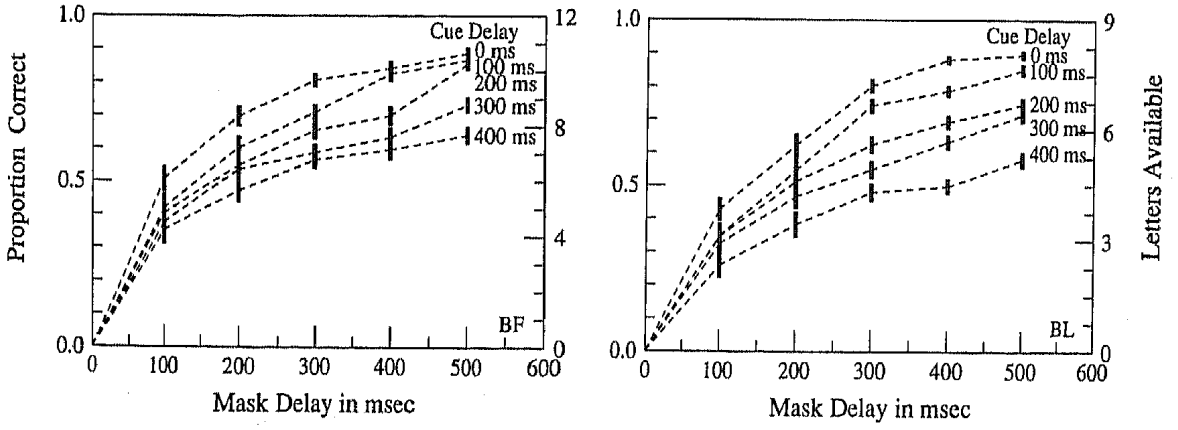

Figure 8. Results of Experiment 2 replotted to show the accuracy of partial reports as a function of mask delay, with cue delay as the parameter. (The vertical bars through the data points indicate the standard error of the proportions. The data points for each cue delay are connected by dotted lines. See Figure 7 for details.)

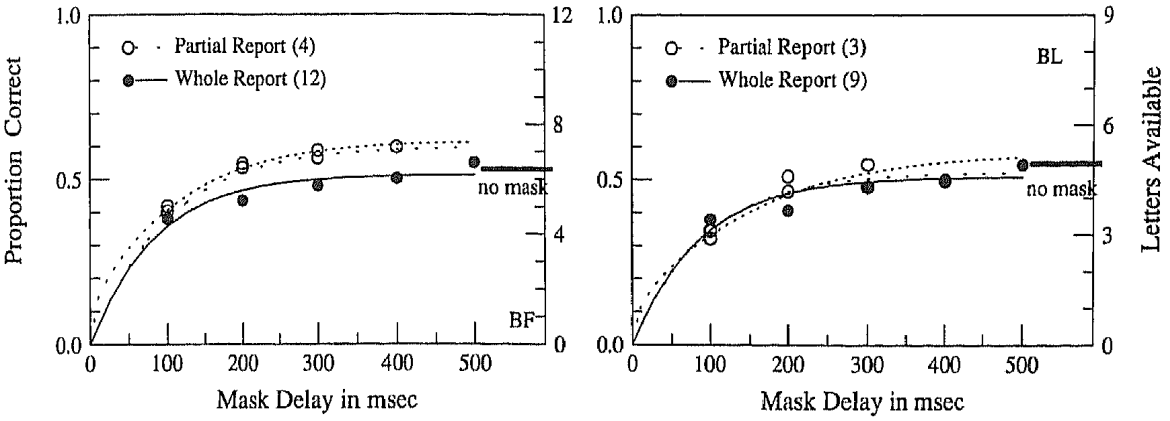

Figure 9. Pure nonselective transfer as a function of mask delay. (The open symbols are the proportion of correct partial reports on trials in which the cue occurred after mask onset [Figure 6b]; the filled symbols are the proportion of correct reports in whole-report-plus-mask trials. The horizontal bar at the right border indicates performance on whole-report trials without masks. The predictions of the aggregate-row model of iconic memory for partial and whole reports are indicated by solid and dotted lines, respectively. The dashed line shows the partial-report predictions of the model described in Equation 13, averaged over the three rows.)

$$L_{c,m} = S_{c,m} + \tfrac{1}{3}N_{c,m}. \quad (1)$$

The factor ⅓ express the fact that because there are three equally likely cues, only one-third of the nonselectively transferred letters are expected to be in the cued row. To apply Equation 1 to our data, we note that we already know two of its three components. If we assume, for the moment, that all letters that are transferred from iconic memory to durable storage are reported, then the partial-report data directly yield the total reported letters, $L_{c,m}$. Partial reports made when the cue occurs simultaneously with or after the mask give the pure nonselective component, $\frac{1}{3}N_{c,m}$ ($c \geq m$), the analysis illustrated in Figure 9. The difference between $L_{c,m}$ and $\frac{1}{3}N_{c,m}$ is the selective transfer, $S_{c,m}$.

Figure 11 shows the values of selective transfer derived from our data. Note that the only difference between Figures 8 and 11 is that the nonselective transfer component has been subtracted from overall performance. All the curves for selective transfer appear to be parallel, shifted vertically. This implies that only one factor determines selective transfer—time elapsed since stimulus onset. Selective transfer that begins, for example, 200 ms after stimulus onset will transfer just as many items in the time period from 200 to 500 ms as selective transfer that began at 0 ms. Because the rate of selective transfer depends only on the elapsed time since stimulus onset, it directly reflects the quality of the stimulus information.

To test the assumption of additivity, we fit the best set of perfectly parallel curves to our data. We do not make any assumptions about the form of the selective or nonselective transfer curves. The solid line segments in Figure 11 all derive from a single curve that has been translated up or down.

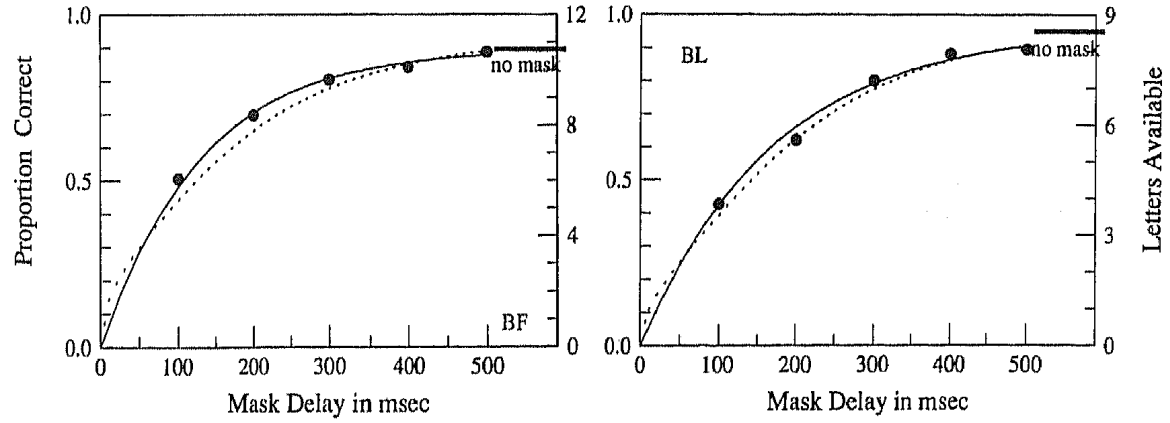

Figure 10. Pure selective transfer as a function of mask delay. (The filled symbols show the proportion of correct partial reports on trials on which the cue occurred at stimulus onset [Figure 6c]. The solid line shows the prediction of the aggregate-row model of iconic memory. The horizontal bar on the right border shows observed performance in a partial-report experiment without masks and a cue delay of 0. The solid line shows the predictions of the model described in Equation 14, averaged over the three rows.)

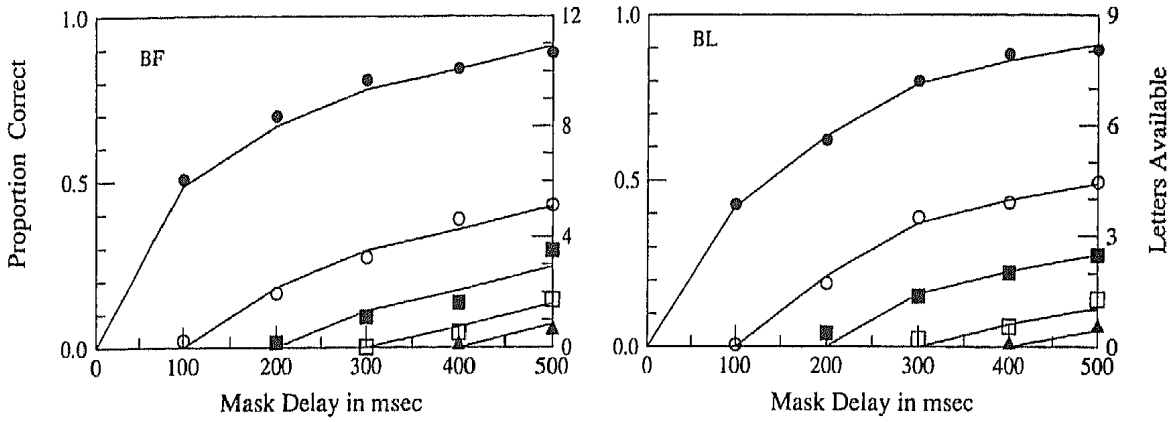

Figure 11. Test of the additivity assumption in the aggregate-row model. (The curves are based on Figure 8—accuracy of partial reports as a function of mask delay, with cue delay as the parameter. Cue delays are in milliseconds: Filled circle = 0, open circle = 100, filled square = 200, open square = 300, triangle = 400. The assumption of algebraic additivity of transfer [Equation 2] permits the subtraction of the estimated nonselective component of transfer [Figure 9] from each curve of Figure 8 to yield the residual selective transfer. The symbols show observed values of residual selective transfer after various cue delays. The solid curves show the predictions that are based on vertical translations of a single generic selective-transfer curve [e.g., delay 0]. The form of the generic selective-transfer curve was estimated from the data.)

The assumption of additivity holds well for our data. The root-mean-square error is 0.016 for subject BL and 0.023 for subject BF.

Some Parametric Assumptions

The pure information transfer functions for the nonselective transfer process in Figure 9 and the selective transfer in Figure 10 can both be approximated by simple exponential growth functions of the form

$$f(t) = C[1 - \exp(-t/\tau)], \tag{2}$$

where C is the asymptotic level of performance, and τ is the time constant of the growth process at an $f(t)$ of 63% of C.

We denote nonselective transfer as $N_{c,m}(t)$, and selective transfer as $S_{c,m}(t)$. The indices remind us that transfer may depend not only on t, but also on the specific values of the cue delay, c, and the mask delay, m. For pure nonselective transfer, the cue comes after the mask, and we obtain

$$N_{c,m}(t) = C_N[1 - \exp(-t/\tau_N)], \qquad c \geq m, \tag{3}$$

where C_N is the capacity of durable storage and τ_N is the time constant for nonselective transfer. Figure 9 shows Equation 3 with the constants chosen to optimize the fit to our data. The deviations of data from theory are very small.

Purely selective transfer occurs when the cue comes at (or before) the stimulus onset. Similarly to Equation 3, we obtain

$$S_{0,m}(t) = C_S[1 - \exp(-t/\tau_S)]. \tag{4}$$

In Equation 4, C_S is the maximum number of letters the subject can transfer from one line. In general, C_S will be very close to the number of letters in one line. However, C_S has to be estimated because subjects are not perfect, and they occasionally miss a letter even in the easiest conditions. The time constant for selective transfer is τ_S.

Figure 10 shows the best fit of Equation 4 to our data for selective transfer. Again, deviations between the theory and the data are small. From Figures 9 and 10, we see that the growth rates of nonselective and selective information transfer curves are quite different, reflecting their different time constants, τ_N and τ_S.

Selective transfer for cue and mask delays with $c \leq m$ occurs only during the interval from c to m:

$$S_{c,m}(t) = S_{0,m}(t) - S_{0,c}(t). \tag{5}$$

Of course, $S_{c,m}(t)$ is 0 whenever $c \geq m$. The total number of letters available for report in the cued row as a function of time is given by a generalization of Equation 1:

$$L_{c,m}(t) = S_{c,m}(t) + \tfrac{1}{3}N_{c,m}(t). \tag{6}$$

The total number of letters available for whole reports is simply $N_{\infty,m}$.

A final complication is that the time a subject needs in order to interpret the cue may be greater than zero. To admit this possibility, a parameter τ_q, the cue interpretation time, is included in the model as an offset parameter, substituting $c + \tau_q$ for c in Equation 6.

Figure 12a summarizes the descriptive model. Two cumulative functions, $N_{c,m}(t)$ and $S_{c,m}(t)$, describe information transfer from iconic memory to durable storage. Before occurrence of a cue, transfer is governed by $N_{c,m}(t)$; after the cue, by $S_{c,m}(t)$. It is useful to think of the cue as a switch that toggles between the two transfer rates.

Predictions for a partial-report-plus-mask experiment are represented in Figure 12b. The number of letters available for report follows the trajectory to C_N until the occurrence of a cue. It then follows the trajectory described by $S_{c,m}(t)$. After the onset of the poststimulus mask, the predicted trajectory is flat.

Figure 12b shows that the cue is predicted to help most

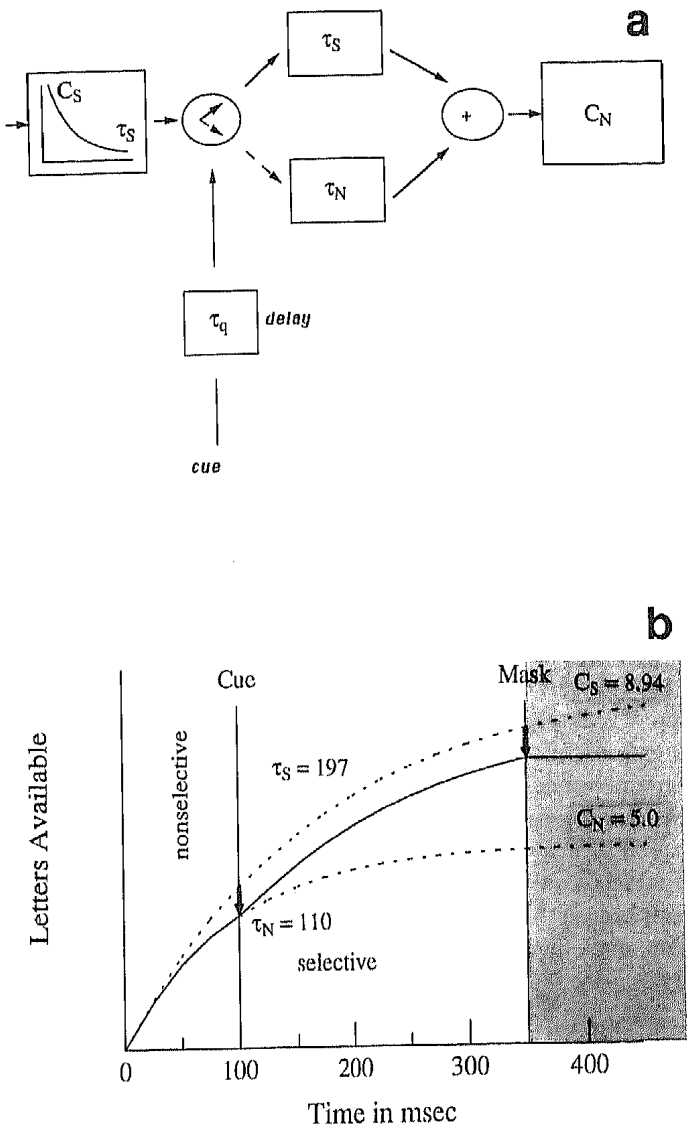

Figure 12. Panel a is a block diagram of the two-process aggregate-row model of iconic transfer. (The first box indicates iconic memory with a capacity C_S decaying with time constant τ_S after a brief stimulus presentation. The partial-report cue, after delay τ_q, causes a shift from an initial nonselective transfer [τ_N] to selective transfer [τ_S] into durable storage. All transferred items are added in durable storage; its apparent capacity C_N varies slightly depending on whether it is determined from partial or from whole reports [see Figure 9].) Panel b illustrates the computation in the two-process model. (Dotted curves show selective and nonselective transfer. Before the cue, transfer is nonselective and proceeds at rate C_N/τ_N to asymptote C_N. After the cue, transfer is selective at rate C_S/τ_S to asymptote C_S. The arrows indicate that in effect, the generic selective transfer curve is joined to the generic nonselective curve at the moment in time that the cue takes effect.)

when given within 100 ms of the stimulus onset. In the first 100 ms, the cumulative transfers $N_{c,m}(t)$ and $S_{c,m}(t)$ differ only slightly. After 100 ms, $N_{c,m}(t)$ reaches its asymptote, whereas $S_{c,m}(t)$ continues for at least another 300 ms.

Parameter Estimations and Fits to the Data

The curves in Figure 7 show the fit of the complete model to the data of both subjects. The model accounts for 96% and 97% of the variance in the data for Subjects BL and BF, respectively. The root-mean-square errors are 0.089 and 0.095, respectively.

The same subjects served in an earlier partial-report experiment, similar to Experiment 1 but without poststimulus masks. These earlier data can be fit by the model derived from Experiment 2 without estimating any new parameters. Figure 13 shows the partial-report-without-masking data and the model predictions. The parameter values were estimated from the experiment using masks. The fit obtained this way does not deviate significantly from the data. The dashed lines in Figure 13 show the contributions of nonselective transfer. The model indicates that, after stimulus termination, selective transfer decreases much faster than one might expect from the relatively slow decay in partial-report superiority.

Table 2 summarizes the parameter estimates for Subjects BF and BL. Two sets of estimates are shown. One set, already described, was derived from the subsets of the data that provided the pure nonselective transfer and the pure selective transfer analyses illustrated in Figures 9 and 10. A second set of parameters was estimated from the complete data of the partial-report-plus-masking experiments. The comparison of these estimates is an indicator of the overall consistency of the model, which is quite good.

For each subject, the nonselective capacity parameters, C_Ns, are very similar in the three relevant data sets: the full partial-report-plus-masking data set, the cue-after-noise subset, and the whole-report data set. The capacities are five letters for BL and seven letters (well above normal) for BF.

The nonselective capacity estimate C_N is effectively equal to 3, the number of letters in one row for BL. It is about 5–10% less than 4 for BF, who was shown four-letter rows.

The time constants for selective and nonselective transfer, τ_S and τ_N, are quite different from each other. Selective transfer continues to rise steadily until after 200 ms, whereas nonselective transfer asymptotes quickly after 100 ms. Both subjects have similar time constants, although their capacities differ.

For both subjects, the time τ_q necessary to interpret the cue is estimated to be 0 or slightly negative.

The speed of the transfer processes is determined by differentiating Equation 2. This results in

$$f'(t) = C/\tau \exp(-t/\tau). \tag{7}$$

For $t = 0$, Equation 7 reduces to

$$f'(0) = C/\tau. \tag{8}$$

Enigmas

Computation of the initial transfer rates immediately at stimulus onset, $S'(0)$ and $N'(0)$, shows that nonselective transfer has a much higher rate. Seventy and 45 letters/s are transferred nonselectively for subjects BF and BL, respectively, and only 27 and 15 are transferred selectively. Note that the nonselective transfer rates are based on the total number of letters transferred into durable storage, not merely on the letters in the cued row. In the aggregate model, it is not obvious why the actual speed of nonselective and se-

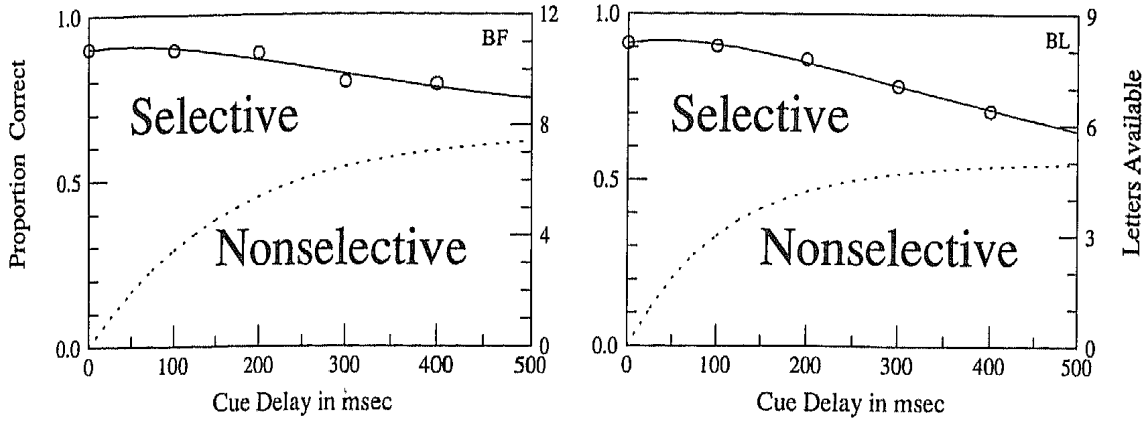

Figure 13. Data from a partial-report experiment without masks. (The open symbols show the proportion correct for various cue delays. The solid line shows the predictions of the two-process aggregate-row model. The dashed curve indicates the estimated component of performance resulting from nonselective transfer. The distance from the dashed line to the solid line [partial-report accuracy] represents the estimated contribution of the selective transfer process.)

lective transfer should appear to differ by so much; this issue is addressed in the row-by-row model, presented later.

It also is surprising that the estimated time the subject needs to interpret the cue τ_q is essentially zero. Experiments by Reeves and Sperling (1986) using visual cues showed that a spatial shift of visual attention took 300–400 ms. Sperling and Weichselgartner (in press) used a click in a go/no-go attention shift experiment that required only turning on attention, not actually shifting it in space. They found a modal switching time of about 100 ms. Our tonal cues, which required a three-choice reaction and a spatial shift of attention, would certainly be expected to have a much longer attention shift latency. These enigmas suggest the need for more complex analysis, which we provide in the next section by analyzing the data separately for each row.

Table 2
Best Fitting Parameter Values for the Aggregate-Row Model (Partial Report With Masking) and for Data Subsets That Yield Estimates of Pure Selective and Pure Nonselective Transfer in Experiment 2

Experiment	C_S	C_N	τ_S	τ_N	t_q	rms error
Subject BF						
Selective	3.58	—	130.0	—	—	0.073
Nonselective	—	7.08	—	99.9	−16.62	0.021
Combination	—	—	—	—	—	0.023
Overall	3.82	7.35	191.1	114.31	−12.5	0.095
Partial report	—	—	—	—	—	0.049
Whole report	—	6.3	—	—	—	—
Subject BL						
Selective	—	—	158.9	—	8.47	0.120
Nonselective	—	4.75	—	104.79	0.29	0.032
Combination	—	—	—	—	—	0.016
Overall	2.98	5.0	197.5	110.4	−13.6	0.089
Partial report	—	—	—	—	—	0.016
Whole report	—	5.0	—	—	—	—

Note. C_S and C_N represent attentional capacities, respectively, of selective and nonselective transfer, with units in letters; τ_s and τ_n represent time constants of selective and nonselective transfer, with units in milliseconds; rms is root mean square. Selective and nonselective experiments the parameters were estimated from subsets of the data that did not require using the additive combination rule. In the combination experiment, the combination rule that estimated only additivity, not any of the parameters, was listed. In the overall experiment, the complete model for Experiment 2 was tested. In the partial-report procedure, the partial-report-plus-mask parameters were used to predict the data from an earlier partial-report-without-mask experiment. The whole-report experiment entailed simply observation of subjects' performance. No parameters were estimated. A dash indicates that a parameter could not be estimated for a particular condition.

Position Effects

Partial-Report Accuracy by Row, Cue Delay, and Mask Onset Time

The probability of correct partial reports as a function of mask delay with cue delay as a parameter is displayed in Figure 14. Each panel shows data for a different row of the display. Partial reports of the middle row differ from reports of the top and bottom rows, and we consider the middle row first. Almost always, subjects report the middle row perfectly. Even in the hardest conditions (short mask delay and long cue delay) subjects report 80% of the middle-row letters correctly. Except for the earliest mask at 100 ms, all the other middle-row curves appear equal at a performance level of about 95% correct. There is no apparent iconic decay for the middle row. Obviously, the transfer to durable storage of letters from the middle row is nonselective, and in this the middle row differs from the other rows.

For the top and bottom rows, performance decreases from near perfect in easy conditions to near chance in the hardest conditions. Because nonselective transfer determines the asymptotic performance at long cue delays, the data for the top and bottom rows indicate there is much less nonselective transfer (and correspondingly more selective transfer) from these rows than from the middle row.

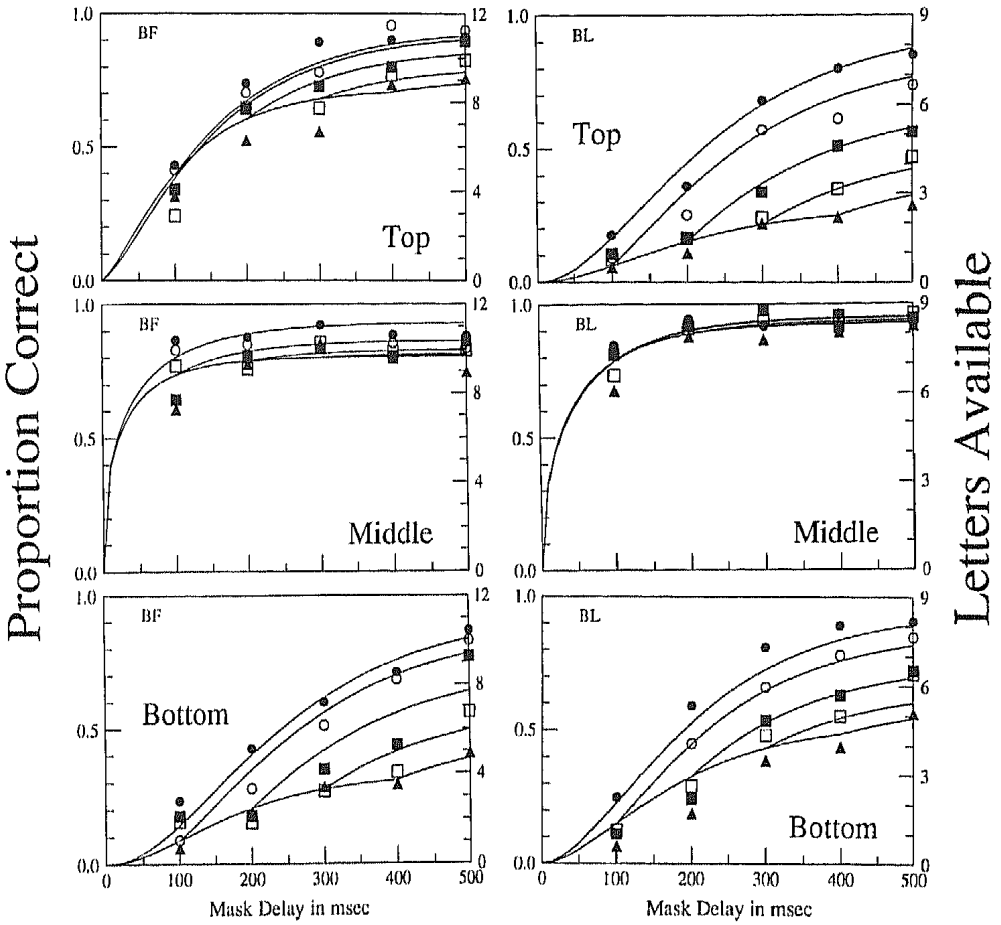

Figure 14. Accuracy of partial reports as a function of mask delay, with cue delay as a parameter, shown separately for each of the three stimulus rows. (Cue delays are in milliseconds: Filled circle = 0, open circle = 100, filled square = 200, open square = 300, triangle = 400. Top, Middle, and Bottom denote the stimulus rows. BL and BF denote the subjects. Each data point represents the proportion correct in 50–100 trials. Note the large and highly significant performance differences between the rows. The curves are predictions of the nine-parameter attentional model [Equations 12–14], with parameters given in Table 3.)

Selective and Nonselective Transfer by Row

To estimate the amount of nonselective transfer, we consider the subset of data with cue onset at or after mask onset (as in Figure 9). Figure 15 shows nonselective transfer for the three rows. For both subjects, the middle row rises to a high asymptotic level within the first 100 ms. The other rows rise slowly and reach generally lower asymptotic levels.

Nonselective transfer: Strategy mixture versus pure strategy. To account for the subjects' good performance on the middle row in the nonselective transfer data, we contrast two possibilities: a trial-to-trial variation of transfer strategy (which, over a series of trials, most often favors the middle row) and a consistent strategy that favors the middle row on every trial. Suppose that, prior to the stimulus exposure on each trial, subjects preselected a particular row to transfer nonselectively immediately following the exposure. Suppose that from trial to trial, they switched their preferred rows, but on the average, they most often chose the middle row. In this strategy, we would expect to find some trials for the top and bottom row on which the subjects' performances were perfect or nearly so. This is not the case, however. Of the trials on which the cue indicated a report of the top or bottom row, fewer than 2% of the reports had all letters correct (compared with 70% for the middle row). From this, we infer that subjects did not switch between rows and that the consistent preference of the middle row in nonselective transfer is responsible for its higher nonselective transfer rate. Indeed, most of the letters reported in the nonselective conditions come from the middle row.

On the other hand, in the conditions that favor selective transfer, the proportions with which the different rows are sampled are nearly the same. This explains why the aggregate model yielded faster rates of letters actually entering durable storage for nonselective than for selective transfer: Nonselective transfer sampled mostly the fast middle row, whereas selective transfer provided an almost equal mixture of all three rows.

Selective transfer. Figure 16 shows pure selective transfer estimated (as in Figure 10) from trials on which the cue occurred simultaneously with the onset of the test flash. For the middle row, selective transfer yields the

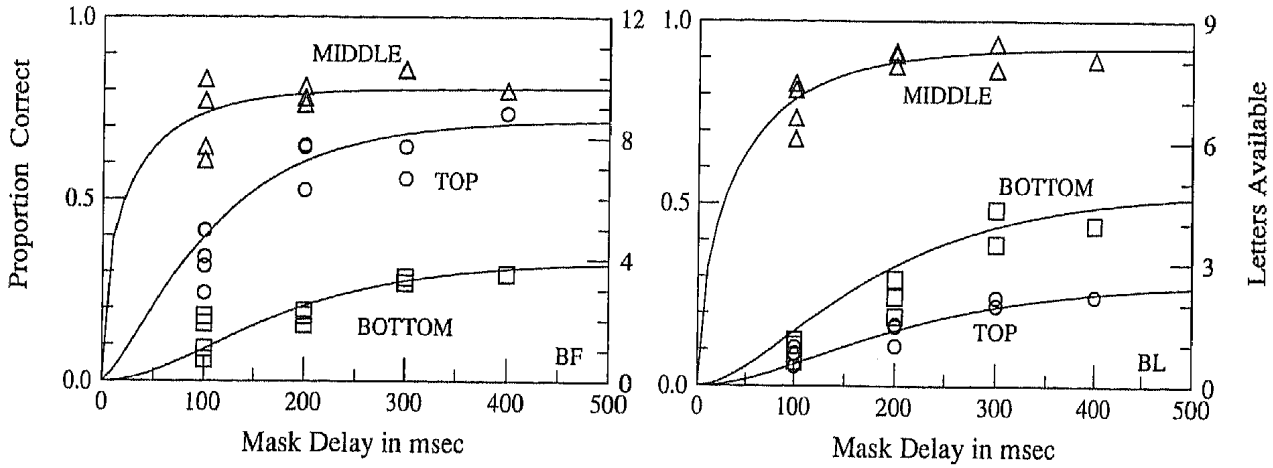

Figure 15. Pure nonselective transfer as a function of mask delay for each of three stimulus rows and 2 subjects. (The data points are the proportion of correct partial reports on trials in which the cue occurred at or after mask onset [Figure 6b]. Circles indicate the top row, triangles indicate the middle row, and squares indicate the bottom row. The solid curves are predictions of the nine-parameter attentional model [Equation 12], with parameters given in Table 3.)

same apparent growth curve as nonselective transfer. Note that, for both subjects, selective transfer for the top and bottom row reaches almost perfect performance at the longest mask delays. This indicates that the information in the stimulus is still available at these long delays. Therefore, it is also available for nonselective transfer. The finding that nonselective transfer reaches asymptote at a lower level for the top and bottom rows must then be a consequence of a capacity limitation. There is a suggestion in the data that the cumulative selective transfer from the top and the bottom rows is an S-shaped function of time. This would mean that the transfer rate was slow in the beginning, reached a maximum value at intermediate times, and finally declined again to zero. A slow start suggests a delayed shift of attention; the slow final rate almost certainly indicates that the iconic image has decayed to illegibility.

The additivity assumption of Equation 1 that yielded selective transfer by subtracting out the nonselective transfer can be applied to the row data to obtain the selective transfer for each individual row. The results are shown in Figure 17. We keep parallelism as a working hypothesis because it accounts for 99% of the variance of the data for both subjects.

Manifestations in the data of attention to a stimulus row. We assume that partial attention to a row slightly improves selective transfer of letters from that row relative to nonselective transfer and that complete attention to a row maximally facilitates selective transfer. Consider a graph of proportion correct versus mask delay with cue delay as the parameter (Figure 14). The earliest mask delay at which data from two cue delays, c_1 and c_2 diverge indicates the point at which the states of attention induced by c_1 and c_2 are sufficiently different to affect selective transfer. For example, consider cues that indicate the bottom row, and suppose $c_1 = 0$ and $c_2 = 100$ ms. In Figure 14, the data for $c_1 = 0$

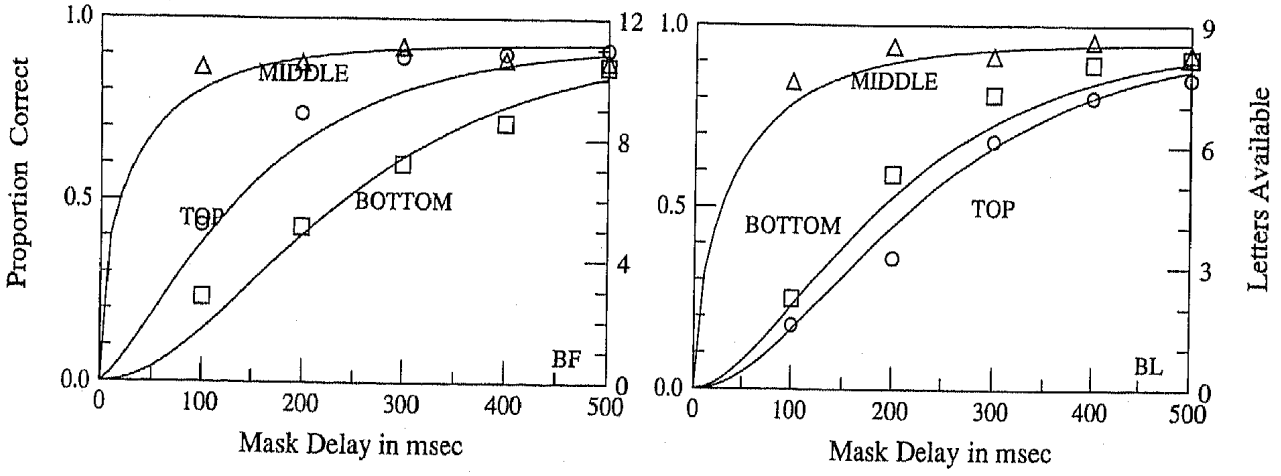

Figure 16. Pure selective transfer as a function of mask delay for each of three stimulus rows and 2 subjects. (Data points are the proportion of correct partial reports on trials in which the cue occurred at stimulus onset [Figure 6c]. Circles indicate the top row, triangles indicate the middle row, and squares indicate the bottom row. The solid curves are predictions of the nine-parameter attentional model [Equation 13], with parameters given in Table 3.)

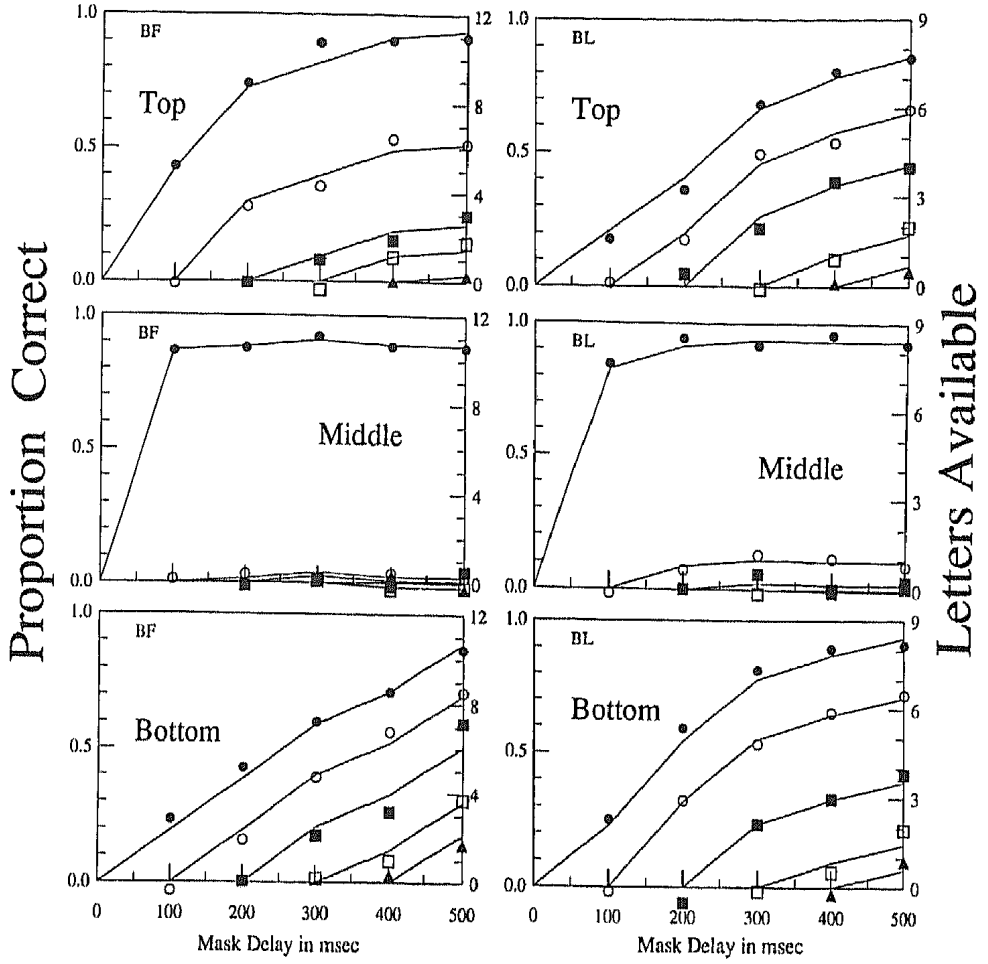

Figure 17. Selective transfer after prior nonselective transfer. (The curves are based on the same data as Figure 14—accuracy of partial reports as a function of mask delay, with cue delay as a parameter. Cue delays are in milliseconds: Filled circle = 0, open circle = 100, filled square = 200, open square = 300, triangle = 400. The estimated amount of nonselective transfer has been subtracted [as in Figure 11] to yield the residual selective transfer. The symbols show estimated values of residual selective transfer after various cue delays.)

first break away from the data for other cs when $m = 100$, and the $c_1 = 0$ data are completely separate when $m = 200$. A mask occurring 100 ms after c_1 means there is no further transfer from the stimulus after 100 ms. For the data obtained with c_1 in this condition to differ from the other c_i implies that the cue must have acted to alter attention within 100 ms. Alternatively, we would have to reject our previous assumption that the mask terminates stimulus availability.

Figure 14 shows that, for the middle row, there is no clear divergence of data for different cue delays and therefore no evidence that attention does or does not affect transfer of the middle row. However, transfer from the top and bottom rows is obviously quite affected by attention. The data for cue delay c in Figure 14 tends to break upward from the pack of longer cue delays as soon as $m \geq c$. This indicates that our cues induce a measurable change in attentional state immediately after their occurrence.

The other aspect of the performance-versus-mask-delay data (Figure 14) that we have already dwelt on at length is the parallelism of the curves for different cue delays onward from the moment $m \geq c$ (Figure 17). Parallelism indicates that the state of selective attention is the same for all the conditions represented in the parallel curve sections. In other words, not only does attention switch quickly once the cue arrives, but it switches completely. If it did not switch all at once, then an early cue, c_1, would have produced a greater attentional shift to the indicated row at a subsequent time, t_2 than a cue, c_2, that did not occur until t_2. In that case, transfer measured at t_2 would be faster for c_1 than for c_2, and the parallelism in the data of Figure 17 would be violated. Because the data are effectively parallel, we also have to assume that within the context of our assumptions, attention shifts quickly and completely. These assumptions are formalized in the next section.

Attentional Model of Transfer From Iconic Memory to Durable Storage

Assumptions

To account for the analysis of partial-report-plus-masking data separately by rows, we generalize the aggregate model of Figure 12a in a natural way, as illustrated in Figure 18. In

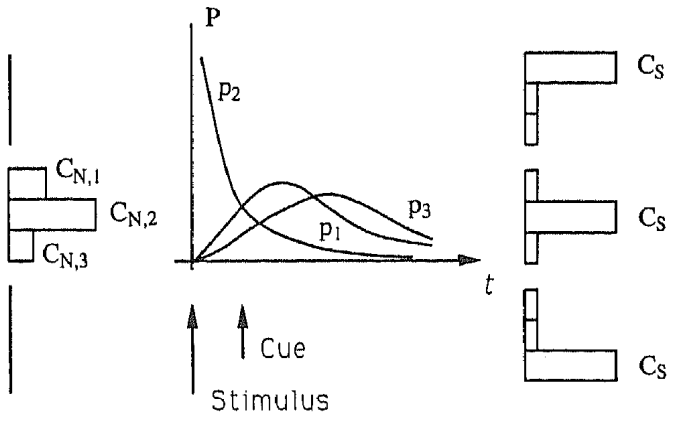

Figure 18. Illustration of the attentional model of iconic memory transfer processes. (As in the aggregate-row model, before the cue, the initial state of attention determines nonselective transfer from iconic memory to a durable store. $C_{N,1}$, $C_{N,2}$, and $C_{N,3}$, respectively, indicate the relative amounts of attention allocated to the top, middle, and bottom rows before the cue. The row-dependent retinotopic component of iconic transfer rate is illustrated by the p_r functions, which begin with stimulus onset. In response to a cue to report Row r, subjects shift attention instantaneously from its initial state to one of the three postcue states indicated by C_S. The actual transfer rate is the product of $C_{N,r}P_r(t)$ before the cue and $C_SP_r(t)$ afterward.)

the aggregate-row model, nonselective transfer and selective transfer were each characterized by a two parameters, their rate and the total capacity. Now these processes are made more explicit in terms of the states of attention they represent. Each state of visual attention is characterized by a spatial function that represents the allocation of attentional resources over space (Sperling & Weichselgartner, 1991) and a separable temporal function that represents the time period during which the spatial function reigns. Thus, nonselective transfer represents the default state of attention that exists from the beginning of the trial until the cue is received and interpreted. The spatial allocation of nonselective attention is described by three numbers ($C_{N,r}$, $r = 1, \ldots 3$) that represent the capacity (in letters) of durable storage allocated to the top, middle, and bottom rows of the display. The single number of the aggregate-row model that described the capacity of durable storage for nonselectively transferred letters was $C_N = \sum_r C_{N,r}$.

After the cue is received and interpreted, one of three states of selective attention occurs. Each state is characterized by C_s, the capacity allocated to the cued row (top, middle, or bottom) and by 0.0 capacity allocated to the other rows (Figure 18).

In the aggregate-row model, the two speeds of transfer from iconic memory to durable storage (nonselective and selective) were parameterized by their own time constants, τ_S and τ_N in Equation 7. In considering rows individually, it is obvious that all transfers are much quicker from the middle row and that three parameters, a_r, $r = 1, \ldots 3$, are needed to characterize the temporal differences between transfer rates from the three different rows. The exact temporal waveforms cannot be determined directly from our data. A mathematically tractable formulation that has useful properties is given by inserting a term $(t/\tau)^{a-1}$ to the time-dependent transfer of Equation 7, resulting in the time-and-row-dependent transfer

$$f'(t, a_r) = C_r(t/\tau)^{a_r - 1} \exp(-t/\tau). \quad (9)$$

The constant C_r is a capacity allocated to row r.

We assume that the differences between nonselective and selective attention are completely captured by the spatial allocations of attentional capacity C_r, so that one set of row-dependent weights, a_r, suffices for all states of attention. Because the a_r depends on spatial location (the row) and does not depend on attention, it represents the intrinsic processing efficiency of a retinal location.

We wish to test our assumption that a single parameter suffices to describe the overall transfer rate of both selective and nonselective attention. Therefore, in parameter estimation, we estimate two overall rate parameters, one for selective and one for nonselective attention, to determine whether these unconstrained rate estimates indeed are similar.

Computational Model

Following Reeves and Sperling's (1986) attentional gating model, it is reasonable to assume that the transfer rate from a location, r, is determined by the product of two factors: (a) the availability (legibility) of stimulus information at r and (b) the amount of attention allocated to r. Availability at a location is determined by iconic buildup and decay, and it is parameterized by exponent, a_r, of Equation 9 combined with the exponential terms. Attentional allocation is parameterized by the capacity allocation, C_r. Equation 9 represents this product. Unfortunately, transfer mode appears implicitly in the time constant, τ, of Equation 9. This means that attention allocation (which determines transfer mode) would be inextricably intertwined with iconic availability if τ_N and τ_S were to differ appreciably.

Cumulative transfer in an interval $[0, m]$ is given by integrating over Equation 9. For simplicity, we change the variable of integration, giving

$$S_{0,m}(m) = \int_0^{m/\tau} C_s(t)^{a_r - 1} \exp(-t)\, dt. \quad (10)$$

By appropriately normalizing the exponential terms in Equation 10, we can convert Equation 10 into an attentional capacity (C, scaled in letters) multiplied by the well-known incomplete gamma function, $P(a, x)$, $0 \le P(a, x) < 1$:

$$P(a, x) = \frac{\int_0^x (t)^{a-1} \exp(-t)\, dt}{\int_0^\infty (t)^{a-1} \exp(-t)\, dt}. \quad (11)$$

In subsequent use, x will take four values: c/τ_N, c/τ_S, m/τ_N, and m/τ_S, representing the intervals from exposure onset to the cue and the mask, respectively, in units of the time constants of nonselective (τ_N) and selective (τ_S) transfer. For $a = 1$, Equation 9 simply reduces to Equation 7. Values of a between 0 and 1 lead to an accelerated exponential growth function for $P(a, x)$. Values higher than

1 lead to S-shaped delayed growth.

Equation 9 plays the same role in the row-by-row model as Equation 7 did in the aggregate-row model. It can be interpreted as representing the time course of the iconic image at each row location weighted by attentional allocation. To help the reader's intuition in following the exposition of the computational model, we show the iconic time course functions for three rows in Figure 18. The instantaneous transfer rate, $p_r(t) = t^{a-1} e^{-t}$, is the derivative of the cumulative transfer, $P(a, x)$, given in Equation 11. As Figure 18 indicates, the rise time for the middle row is too fast to be observed in our experimental conditions; however, the middle row decays with what appears to be a familiar, monotonically decreasing function. The top and bottom rows rise before they decay. We defer to later the question of whether these top- and bottom-row functions truly represent the rise and decay of iconic memory. (Alternatively, they might represent a property of a model in which the absence of independent measurements of attention insufficiently constrains the partition of transfer rates into legibility and attentional components.)

The following equations summarize the model. Equation 12 describes the cumulative nonselective transfer that takes place from the onset of the stimulus until either a cue or a mask occurs. It is the product of terms representing two factors: attentional allocation, C, and retinotopic/stimulus factors, P:

$$N_{r,c,m} = C_{N,r} P(a_r, c'/\tau_N), \qquad c' = \min(c, m). \quad (12)$$

The function P implies different transfer dynamics for each of the three rows r. Relative to the middle row, transfer from the top and middle rows is both delayed and slower. Because delay and slowness are perfectly correlated, both are captured by the parameter a_r.

There is a limit to the total number of letters that can be transferred from iconic to durable storage. How the subjects allocated space in durable storage to particular rows so as to optimize their performance is a matter that we did not attempt to control. Therefore, a parameter $C_{N,r}$ is needed for each row to describe the maximum number of letters of durable storage allotted to it (i.e., the default allocation of attention prior to the cue). Finally, the overall rate of nonselective transfer is determined by the time parameter τ_N.

Equation 13 describes selective transfer. In this formulation, selective transfer begins instantly at the onset of the cue and ends instantly at the onset of the mask:

$$S_{r,c,m} = C_S[P(a_r, m/\tau_S) - P(a_r, c/\tau_S)], \qquad m > c. \quad (13)$$

The cumulative transfer to durable storage depends on the integrated product of available information and attention (Reeves & Sperling, 1986). Available information is represented here by $P(a, x)$, which, for the special case of $\tau_N = \tau_S$ depends only on elapsed time since onset of the stimulus. Attention is represented by the currently operative set of C_r capacity values. Attention depends only on elapsed time since the onset of the cue. Therefore, when $\tau_N = \tau_S$, iconic time course and attention are independent.

Equation 14 expresses that the total number of letters transferred to durable storage from each row r is the sum of the nonselective and selectively transferred letters from r. It generalizes Equation 1 of the aggregate model:

$$L_{r,c,m} = N_{r,c,m} + S_{r,c,m}. \quad (14)$$

Parameter Estimates and Their Interpretation

Best fitting parameters were estimated for Equations 12–14 by means of an optimization program (PRAXIS; see Brent, 1973; Gegenfurtner, 1992), using the partial-report-plus-masking data of Experiment 2. The results of parameter estimation are summarized in Figure 14 and Table 3. The model's predictions correlate very well with the data: $r^2 = .98$ and .95 for the 2 subjects. The predicted average selective and nonselective transfer for the three rows is almost identical to the predictions of the aggregate-row model (see Figures 9 and 10). Therefore the row-by-row model (without additional parameters) also predicts the data from the whole-report experiment and from the partial-report experiment without masks.

The time constants τ_S for selective transfer and τ_N for nonselective transfer are now both approximately 100 ms, indicating that each transfer process completes in about the same time. However, the actual transfer rates C_r/τ depend on the row capacity. The aggregate-row model's capacity for nonselective transfer, C_N, is now split up into the three $C_{N,r}$s. This set of rates defines the initial default attention state prior to the cue. The high rate for the middle row indicates that the default attention state is primarily focused on the middle row. When a cue is received and interpreted, attention is shifted to the cued row, and the transfer rate is determined by the iconic legibility of that row, $f'(t, a_r)$ (Equation 9). In fact, the selective capacity, C_S, is virtually the same as in the aggregate-row model and nearly equal to the number of letters in the row. In effect, the model assumes that once attention is shifted away from the center row to the top or bottom row, it is as effective at the top or bottom row as it was in the center, and any difference in performance must be accounted for by differences in iconic legibility.

Parallel versus serial process in nonselective transfer.

Table 3
Best Fitting Parameter Values for the Model That Takes Differences Between Rows Into Account

Row	C_S	C_N	τ_S	τ_N	a
		Subject BF			
Top	3.72	2.87	115	82	1.39
Middle	3.72	3.21	115	82	0.38
Bottom	3.72	1.31	115	82	2.27
		Subject BL			
Top	2.85	0.85	109	97	2.25
Middle	2.85	2.78	109	97	0.50
Bottom	2.85	1.59	109	97	1.97

Note. C_S and C_N represent attentional capacities, respectively, of selective and nonselective transfer, with units in letters; τ_s and τ_n represent time constants of selective and nonselective transfer, with units in milliseconds; a is a pure number (Equation 12) that represents attentional dynamics.

The a_r parameters represent the dynamics of buildup and decay of iconic legibility at the retinal locations r. They represent the availability of information from a particular row regardless of whether it has been cued. In fact, some time after stimulus termination (100 ms for Subject BF, 200 ms for Subject BL; see Figures 15 and 17) the slopes of the nonselective transfer functions for the top and bottom rows in Figure 15 are still as steep or again become as steep as the initial slopes. This means that after 100 or 200 ms, the availability of information from these rows is as well as or better than it is immediately after stimulus termination.

One interpretation of the delayed availability of information from the top and bottom rows in nonselective transfer is that iconic legibility builds up slowly but approximately simultaneously in these noncentral rows. An alternative explanation is based on serial processes. Prior to a cue, subjects preprogram their attention to move away from fixation at about the time they expect to have completed transfer of the middle row to durable storage. Then they shift attention randomly to either the top or bottom row. This would result in an apparent delay in the availability of information from the top and bottom rows. If subjects were indeed shifting attention on nonselective report trials, it would greatly complicate the analysis of the attentive and iconic components of performance. The present data do not discriminate well between these alternatives.

Attention and the iconic time course are inextricably bound by multiplication in Equation 9: Only the product of attentional allocation and iconic availability determines performance. The model is a powerful computational device, but without an independent verification of the attentional state (or iconic availability), it is not a sufficiently precise tool to dissect unambiguously the attentional and iconic components of performance.

A comparison of the estimated values of a_1 and a_3 in Table 3 shows that the 2 subjects' iconic time courses are different in the top and bottom rows, with BF favoring the top row and BL the bottom row. Because these effects occur in both nonselective and selective attention, the model assigns them to the iconic time course. However, a more plausible interpretation would suggest that they represent tendencies, or biases, to shift attention up or down. According to this interpretation, in response to a cue, BF shifts his attention faster to the top row than the bottom row, and for BL it is just the reverse. These biases in selective transfer mirror the subjects' bias in nonselective transfer.

Although the model assumes that nonselective transfer reflects a single initial attentional state, closer examination of the data suggests that subjects first transfer letters from the middle row and then fill up the remainder of durable storage using nonselective transfer from the other rows. That is, even nonselective transfer ultimately may have to be modeled as consisting of two or more attentional states, an initial state of attention to the middle row, followed by attention to either the top or bottom row. Although this precision of description is necessary for the accurate partition of the components of performance (iconic decay, attention), it is not necessary from a purely computational point of view. The model accounts nicely for all the enigmas that remained after the aggregate-row model and provides a framework for dealing with the few problems that remain.

General Discussion

The present data show the critical importance of nonselective transfer in iconic memory experiments. By decomposing performance into selective and nonselective transfer and subtracting nonselective transfer from the total transfer, we were able to isolate the selective component that depends on the stimulus decay and attentional shifts. This isolation of the two transfer processes was made possible by using a completely crossed design of cue delay and mask delay. This crossed design differs from previous investigations with poststimulus masks (e.g., Averbach & Coriell, 1961; Irwin & Brown, 1987), in which only one mask or cue delay was used or cue and mask delay were correlated.

With respect to theory, we consider the three previous computational treatments of information transfer from iconic memory to durable storage. The earliest model (Averbach & Coriell, 1961) is extremely simple because it was developed for a more restricted paradigm. It proposes both a selective and a nonselective transfer process, but it embodies an assumption about probabilistic independence between these two processes that is strongly contradicted in our larger data set. Rumelhart's (1970) model is quite similar to ours. It fails because it embodies an incorrect assumption about subjects' strategies and another about memory capacity limits. These two models, and ours, share the common theme of two transfer processes. The third model (Loftus et al., 1985) derives iconic decay properties from a single nonselective transfer process. With respect to nonselective transfer and iconic decay, there is considerable agreement between Loftus et al.'s theory and ours, although their theory is not intended to confront the two transfer process issues that are our primary concern. In the next three sections, we consider these models in more detail. Then we briefly review noncomputational suggestions about iconic transfer processes.

Probabilistic Independence of Nonselective and Selective Transfer

Averbach and Coriell (1961) did a partial-report experiment in which the stimulus was two rows of 8 letters and the required partial report was a single letter. A visual cue ("bar marker") appeared above or below the required letter. Total transfer was determined in a partial-report experiment with a cue to report 1 of 16 possible letters. Nonselective transfer was estimated from report accuracy when the cued letter was masked with a concentric annulus. Selective transfer was estimated by correcting total transfer for the nonselective component. Because Averbach and Coriell's annulus was an effective letter masker only when the annulus occurred after a letter, and not when it occurred simultaneously, they ignored the data of the initial

parts of their masking curves. However, their experimental results are generally similar to ours, even though the experimental conditions are quite different. Overall, performance was higher for our subjects.

Averbach and Coriell (1961) proposed the following combination rule. Their basic unit of analysis was a single letter, which could be transferred either selectively or nonselectively. They regarded the two transfer types as independent processes. Both transfers contribute probabilistically to the proportion of correctly reported letters, much as in Rumelhart's (1970) model. Averbach and Coriell found huge performance differences for different letter positions but decided to ignore them and average their data. Moreover, they did not vary mask and cue delays independently, so they were severely limited in what they were able to do with their data and theory. For example, they were noncommittal about whether nonselective transfer ends when the cue occurs, about whether selective transfer begins immediately upon cue onset, and about other issues related to the underlying processes.

Averbach and Coriell's (1961) model is analyzed as follows. Each letter has a certain probability of being transferred by either process. Denote the event of a nonselective transfer with N, the event of a selective transfer with S, and the event of any kind of transfer with T:

$$P(T) = P(N) + (1 - P(N))P(S). \quad (15)$$

It then follows that selective transfer is given by

$$P(S) = (P(T) - P(N))/(1 - P(N)). \quad (16)$$

Equation 16 expresses the idea that two processes contribute to partial-report accuracy, as does our Equation 2 (and its subsequent elaborations) and Rumelhart's (1970) model (discussed shortly).

Figure 19 compares the rates of selective transfer (i.e., the legibility of the iconic image) as derived from the present model and from Averach and Coriell's (1961) model. It shows the number of letters selectively transferred during successive 100-ms intervals plotted as a function of the time at the end of the interval. The data points in Figure 19a are derived from Equation 2. Figure 11 shows cumulative selective transfer; Figure 19 shows selective transfer rate, that is, successive differences between the points of the lines in Figure 11. That all these successive differences fall on the same iconic decay function should be no surprise. We previously noted that all the curves of Figure 11 derive from a single generic selective transfer function.

Figure 19b shows the predictions for Averbach and Coriell's (1961) formulation (assuming that nonselective transfer stops after the occurrence of the cue). When the cue delay is zero, there is no nonselective transfer, and both our model and theirs give the same predictions (indicated by the filled circles). However, when nonselective and selective transfer are combined (i.e., for any cue delay greater than zero), the models differ. As already pointed out in the discussion of Figure 11, our assumptions result in selective transfer rates that depend only on the time since the onset of the cue. Averbach and Coriell's model leads to large, highly irregular estimates of selective transfer for a given time interval, and no clear pattern emerges of how selective transfer is determined. Their model cannot account for our data.

The consistency of the different independent estimates of the iconic decay function demonstrates the value of our Cue Delay × Masking Delay crossed design, which enables us not only to estimate the parameters for our model, but also to check our model's consistency.

Diffuse Transfer Followed by Focused Transfer

Rumelhart (1970) proposed a mathematical model of partial-report experiments cast in terms of features. Features were transferred with replacement from retinal locations and aggregated to form letters. During the stimulus exposure, features were equally available at all locations and all times. After termination of the exposure, feature availability was assumed to decay exponentially. The feature extraction rate was assumed to have an absolute limit (capacity). Before a cue was received, the overall feature extraction capacity was spread equally over all locations. Immediately after a cue was received, feature extraction capacity was concentrated entirely on the cued locations.

The essential ideas of Rumelhart's (1970) model are quite similar to those of our model, namely, that there is a default precue attentional state followed by a postcue attentional state and that the same transfer process operates in both states (merely the row allocations are different). However, Rumelhart was unaware that the precue state is not diffusely spread over all rows but is concentrated on the middle row. In addition, he had no explicit capacity limit for durable storage, relying on limited stimulus availability to account for all response limitations. This was obviously too restrictive an assumption.

Rumelhart's (1970) representation of the probability of correct reports, P, as the indirect result of a feature extraction process would allow the P versus time graphs either to grow like exponentially limited growth processes or to assume S shapes. An S shape would result from the fact that before a threshold number of features is collected at a location, the probability of correctly reporting the letter at that location is assumed to be at chance. After the critical number of features is collected, the probability of correct report is assumed to be 1.0. Although there is considerable flexibility in the generation of S-shaped curves under the feature accumulation assumption, and our empirical P versus t curves are, in a few cases, S shaped, it seemed better not to burden our transfer theory with such a complex assumption.

How Much Is an Icon Worth?

The nonselective transfer curves obtained in our experiments appear very similar to the ones derived by Loftus et al. (1985) in a paradigm using pictorial stimuli. They measured the number of details subjects could report from briefly exposed pictures. Exposure duration was varied, and a mask followed stimulus presentation immediately after stimulus offset. In a second condition, presentation of the mask was delayed 300 ms. They found that a 300-ms mask delay after

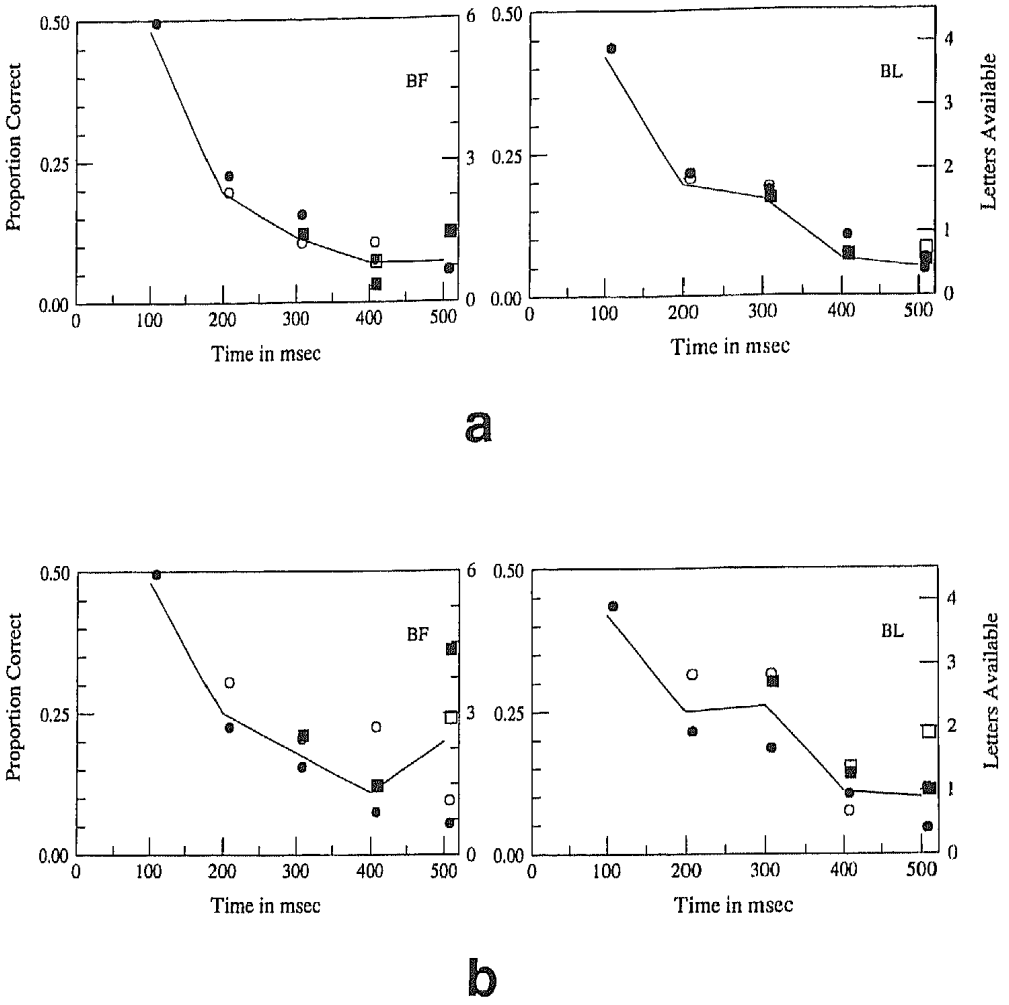

Figure 19. Derived iconic memory decay functions: Estimates of the rate of selective transfer at a given time after onset of the 50-ms stimulus. (At each time *t*, selective transfer during the 100 ms preceding *t* is estimated independently from each condition, with cue delay less than *t*. [This requires extraction of the selective transfer component from total transfer whenever cue delays are greater than 0.] Symbols indicate cue delays in milliseconds: Filled circle = 0, open circle = 100, filled square = 200, open square = 300, triangle = 400. Data are shown for Subjects BF [left] and BL [right]. Panel a shows estimates of selective transfer derived from our model [Equation 5]. Data points are the differences between successive points on each of the lines of Figure 11. Insofar as the different estimates all fall on the same iconic decay function, it substantiates our model of iconic decay. Panel b shows estimates of selective transfer derived from Averbach and Coriell's [1961] model [Equation 16] applied to our data. The wide variation [at a given time] of the different estimates for selective transfer indicates that this model does not yield a consistent description of iconic decay.)

the termination of a stimulus exposure led to the same levels of performance as an additional 100-ms stimulus exposure. They argued that an additional exposure of 100 ms is equivalent to an icon that is available for 300 ms. In this and subsequent experiments (Loftus, Duncan, & Gehrig, 1992; Loftus & Hogden, 1988), with various stimulus materials and tasks, they equilibrated iconic availability against an equivalent continued exposure that yielded the same performance. Ultimately, Loftus et al. (1992) derived an iconic decay function in terms of equivalent continuation of the stimulus. Their derived iconic decay functions were approximately, but not precisely, exponential.

The assumptions underlying Loftus et al.'s (1992) and our analyses are quite similar, although they derived all their data from whole reports. The main difference is that we present iconic decay directly in terms of a transfer rate; whereas they presented it in terms of the fraction of the transfer rate of a continued stimulus exposure. Furthermore, in their procedures, it apparently was not necessary to discriminate the transfer rate at different retinal locations, which is critical in our analyses. Their derived iconic decay functions agree quite well with those we derive for the middle row.

Position Effects

Differences in performance for different parts of the display have long been observed (Averbach & Coriell, 1961; Holding, 1970; Sperling, 1960); but have not been taken into

account in the estimation of the duration of iconic memory. The differences we observe are mainly lower transfer rates for the top and bottoms compared with the middle row. Although in our formulation these locational factors are tied to the stimulus, it is likely that they are based, at least in part, on attentional factors. A relevant positional analysis was done by Holding (1970). He varied the probability with which each row was cued and found that performance varied accordingly, implicating attention. However, Holding's analysis was insufficient to discriminate a change in precue nonselective strategy from (postcue) difference in iconic decay. Furthermore, like many others (see Long, 1980), we strongly disagree with Holding's conclusion that this observation can explain partial-report superiority without postulating an intermediate store.

When a rapidly moving spot is illuminated by stroboscopic flashes (temporally sampled motion), more than one spot appears to move simultaneously. The number of apparently visible spots is a measure of visible persistence, and this number varies with retinal location, persistence being longer for peripheral than for foveal stimuli (Farrell, Pavel, & Sperling, 1990). Eventually, such spatial nonhomogeneities of the visual system will have to be reflected in accounts of iconic decay.

Strategy

The results obtained in Experiment 1 seem to contradict earlier results by Sperling (1960) that showed an influence of subjects' strategy. In Experiment 1, performance for a given cue delay varied depending on which cue delays were given in preceding sessions. We can resolve this contradiction by looking at subjects' overall performance level. Our subjects were well practiced. Under ideal conditions (no mask or cue delay), they achieved a performance level of 95–100% correct. Subjects in Sperling's study achieved 70–95% correct under the same conditions. This suggests that these early strategy effects were due to the use of nonoptimal strategies that would have been discarded after additional practice.

Other Models

In recent studies, Irwin and Brown (1987) and Irwin and Yeomans (1986) tested an alternative conception of iconic memory. Their theory also has two buffers, analogous to the iconic memory and short-term memory concepts of traditional theories. It assumes that the coding in iconic memory has two separate representations, one for identity and another for spatial location. This distinction does not bear directly on the distinction between nonselective and selective transfer, but it is relevant to the general issue.

We did not analyze our data for errors of intrusion and location, but we have some important relevant observations from serving as subjects ourselves and from speaking to subjects in iconic memory experiments. When a subject happens to be attending to one row while another one is cued (and fails to perceive the cued letters), the tendency is not to guess at random, but to report the (nonselectively) transferred letters even though they are known to be in the wrong row. To a subject, it seems better to report letters at least known to have been somewhere in the stimulus than to report random letters; the reasoning perhaps being that the cue or the rows may have been misperceived. Therefore, in assessing location errors, it is critical to use additional measures to assess the nature of the errors. For example, Sperling and Dosher (1986) noted that when items were reported with high confidence, location errors were extremely rare and practically never extended beyond an adjacent location. Irwin and Yeomans (1986) supported the notion of row juxtaposition. In their study, incorrect letters mostly came from an incorrect row in the correct column of the display.

Summary and Conclusion

We experimentally identified two transfer processes, nonselective and selective, in the partial-report task. Our data provided strong evidence that performance in the partial-report task is given by the algebraic sum of these two processes. Experiment 1 showed that independent of cue delay, subjects use only one strategy in a partial-report experiment. Experiment 2 showed that this strategy consists of nonselectively transferring letters until the cue appears and afterwards selectively transferring them.

The many complexities of these experiments are accurately described by a computational model that makes several plausible assumptions. Transfer rates are determined by the product of iconic legibility of the stimulus (which depends on the elapsed time after stimulus exposure and on the retinal location) and the subject's attentional state. Nonselective transfer is characterized by rapid transfer of the middle row and much slower transfer of other rows. This precue attentional state is parameterized in the computational model by the precue capacity allocations weighted heavily toward the middle row.

Immediately after the cue, attention shifts to the cued row of the display. Postcue capacity allocation is maximum for the cued row and zero for the others. From this moment on, until the poststimulus mask ends all iconic transfer, selective transfer occurs from the cued row. Nonselective transfer is focused mainly on the middle row, whereas selective transfer focuses exclusively on the cued row, so that selective transfer produces more correct items on the average—a higher effective transfer rate. However, empirically determined rate constants (completion times) for nonselective and selective transfers are approximately the same ($\tau \approx 100$ ms), suggesting that all transfers represent the same process and the different effective rates reflect different states of attention, different retinal locations, and different likelihoods that the transferred items will be in the cued row.

References

Adelson, E. H., & Jonides, J. (1980). The psychophysics of iconic storage. *Journal of Experimental Psychology: Human Perception and Performance, 6,* 486–493.

Averbach, E., & Coriell, E. (1961). Short-term memory in vision.

Bell System Technical Journal, 40, 309–328.

Averbach, E., & Sperling, G. (1960). Short term storage of information in vision. In C. Cherry (Ed.), *Information theory* (pp. 196–211). London: Butterworth.

Brent, R. P. (1973). *Algorithms for minimization without derivatives.* Englewood Cliffs, NJ: Prentice Hall.

Budiansky, J., & Sperling, G. (1969). *GSLetters. A general purpose system for producing visual displays in real time and for running psychological experiments of the DDP24 computer* (Bell Telephone Laboratories Technical Memorandum 69-1223-6). Murray Hill, NJ: AT&T Bell Laboratories.

Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & Psychophysics, 27,* 183–228.

Dick, A. O. (1969). Relations between the sensory register and short-term storage in tachistoscopic recognition. *Journal of Experimental Psychology, 82,* 279–284.

DiLollo, V. (1984). On the relationship between stimulus intensity and duration of visible persistence. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 144–151.

Duncan, J. (1983). Perceptual selection based on alphanumeric class: Evidence from partial reports. *Perception & Psychophysics, 33,* 533–547.

Efron, R. (1970). Effect of stimulus duration on perceptual onset and offset latencies. *Perception & Psychophysics, 8,* 231–234.

Eriksen, C. W., & Collins, J. F. (1967). Some temporal characteristics of visual pattern perception. *Journal of Experimental Psychology, 74,* 476–484.

Farrell, J. E., Pavel, M., & Sperling, G. (1990). The visible persistence of stimuli in stroboscopic motion. *Vision Research, 30,* 921–936.

Gegenfurtner, K. R. (1992). PRAXIS: Brent's algorithm for function minimization. *Behavior Research Methods, Instruments, & Computers, 24,* 560–564.

Hall, D. C. (1974). Eye movements in scanning iconic imagery. *Journal of Experimental Psychology, 103,* 825–830.

Hogden, J. H., & Di Lollo, V. (1974). Perceptual integration and perceptual segregation of brief visual stimuli. *Vision Research, 14,* 1059–1069.

Holding, D. H. (1970). Guessing behavior and the Sperling store. *Quarterly Journal of Experimental Psychology, 22,* 248–256.

Irwin, D. E., & Brown, J. S. (1987). Tests of a model of informational persistence. *Canadian Journal of Psychology, 41,* 317–338.

Irwin, D. E., & Yeomans, J. M. (1986). Sensory registration and informational persistence. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 343–360.

Kaufman, J. (1978). *Visual repetition detection.* Unpublished doctoral dissertation, New York University, New York.

Kahneman, D. (1968). Method, findings, and theory in studies of visual masking. *Psychological Bulletin, 70,* 404–425.

Kropfl, W. J. (1975). *Variable raster and vector display processor* (Bell Telephone Laboratories Technical Memorandum 75-1223-3). Murray Hill, NJ: AT&T Bell Laboratories.

Külpe, O. (1904). Versuche über Abstraktion [Experiments on abstraction]. In V. F. Schumann (Ed.), *Bericht über den 7. Kongress für experimentelle Psychologie* (pp. 56–68). Leipzig: Barth.

Loftus, G. R., Duncan, J., & Gehrig, P. (1992). On the time course of perceptual information that results from a brief visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 530–549.

Loftus, G. R., & Hogden, J. (1985). Extraction of information from complex visual stimuli: Memory performance and phenomenological appearance. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 139–191). San Diego, CA: Academic Press.

Loftus, G. R., Johnson, C. A., & Shimamura, A. P. (1985). How much is an icon worth? *Journal of Experimental Psychology: Human Perception and Performance, 11,* 1–13.

Long, G. M. (1980). Iconic memory: A review and critique of the study of short-term visual storage. *Psychological Bulletin, 88,* 785–820.

Melchner, M. J., & Sperling, G. (1980). *VEX: A computer system for real-time vision experiments* (Bell Telephone Laboratories Technical Memorandum 80-1223-4). Murray Hill, NJ: AT&T Bell Laboratories.

Merikle, P. M., Lowe, D. G., & Coltheart, M. (1971). Familiarity and method of report as determinants of tachistoscopic performance. *Canadian Journal of Psychology, 25,* 167–174.

Mewhort, D. J. K., Campbell, A. J., Marchetti, F. M., & Campbell, J. I. D. (1981). Identification, localization, and "iconic memory": An evaluation of the bar-probe task. *Memory & Cognition, 9,* 50–67.

Mewhort, D. J. K., Johns, E. E., & Coble, S. (1991). Early and late selection in partial report: Evidence from degraded displays. *Perception & Psychophysics, 50,* 258–266.

Mewhort, D. J. K., Merikle, P. M., & Bryden, M. P. (1969). On the transfer from iconic to short-term memory. *Journal of Experimental Psychology, 81,* 89–94.

Neisser, U. (1967). *Cognitive psychology.* Englewood Cliffs, NJ: Prentice Hall.

Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. In H. J. Pick & I. J. Saltzman (Eds.), *Modes of perception* (pp. 137–157). Hillsdale, NJ: Erlbaum.

Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review, 93,* 180–206.

Rumelhart, D. E. (1970). A multicomponent theory of the perception of briefly exposed visual displays. *Journal of Mathematical Psychology, 7,* 191–216.

Sakitt, B. (1976). Iconic memory. *Psychological Review, 83,* 257–276.

Scarborough, D. L. (1972). Memory for brief visual displays of symbols. *Cognitive Psychology, 3,* 408–429.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74,* 1–29.

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

Sperling, G. (1967). Successive approximations to a model for short-term memory. *Acta Psychologica, 27,* 285–292.

Sperling, G., & Dosher, B. (1986). Strategy and optimization in human information processing. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and performance* (Vol I, pp. 1–65). New York: Wiley.

Sperling, G., & Gegenfurtner, K. (1988). Two transfer processes in iconic memory. *Bulletin of the Psychonomic Society, 26,* 488.

Sperling, G., & Weichselgartner, E. (in press). Episodic theory of the dynamics of spatial attention.

Townsend, V. M. (1973). Loss of spatial and identity information following a tachistoscopic exposure. *Journal of Experimental Psychology, 98,* 113–118.

Wundt, W. (1899). Zur Kritik tachistoskopischer Versuche [A criticism of tachistoscopic experiments]. *Philosophische Studien, 15,* 287–317.

Received March 28, 1991
Revision received September 30, 1992
Accepted October 6, 1992 ■