THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY
2006, 59 (1), 150–160

Psychology Press
Taylor & Francis Group

Do masks terminate the icon?

Hannah Smithson and John Mollon
University of Cambridge, Cambridge, UK

Iconic memory is operationally defined by part-report experiments (Sperling, 1960). If a mask is presented after the target, the mask is thought to be superposed on the target in the iconic representation, or to displace it from the representation. But could a cue presented *after* a pattern mask still allow selection within the target array? A target array of letters was followed by a checkerboard mask. We compared two target–mask interstimulus intervals (ISIs; 0 and 100 ms), and six cue delays. At ISI = 0 ms, performance was at chance, for part report and whole report. At ISI = 100 ms, with the shortest cue delay, observers demonstrated a part-report advantage of 25–30%. As cue delay increased the part-report advantage decreased. These results are inconsistent with an iconic memory that is automatically displaced or overwritten by new information. We consider two alternatives: a second-stage store, which represents letters in terms of their high-level features and which the mask cannot penetrate, or a four-dimensional store that preserves separately the representations of the target and its aftercoming mask. We discuss the implications of our results for studies that use backward masking to "terminate the icon".

Real-life tasks, such as listening to speech, reading a report, or riding a bicycle, require synthesis of sensory information over time. It is believed that the brain has the capacity to store, for a short time, a large amount of sensory information, in a raw form (Averbach & Coriell, 1961; Chow, 1986; Coltheart, 1983; Eriksen & Collins, 1967; Neisser, 1967; Palmer, 1988; Sperling, 1960, 1963; Treisman, Russell, & Green, 1975). In the case of vision, this storage is thought to survive for no more than a few hundred milliseconds and to be vulnerable to an aftercoming mask. It has been distinguished from a more durable short-term visual memory that survives for many seconds, has lower capacity, and is not degraded by a pattern mask that is irrelevant to the task (Phillips, 1974; Phillips & Baddeley, 1971). In the present paper we investigate the way in which very-short-term visual storage preserves successive inputs.

In many studies of attention, short-term memory, and language, a patterned mask is presented to control the amount of time that a target stimulus is available for perceptual

---

Correspondence should be addressed to Hannah Smithson, Department of Psychology, Durham University, Science Laboratories, South Road, Durham, DH1 3LE, UK. Email: hannah.smithson@durham.ac.uk

Hannah Smithson is now at the Institute of Ophthalmology, University College London.

Note: This material was presented at the London meeting of the Experimental Psychology Society in January, 2004.

We would like to thank Dr. R. Sorensen for acting as an observer, and Prof. G. Sperling and Prof. K. Gegenfurtner for helpful discussion.

http://www.psypress.com/qjep

DOI:10.1080/17470210500269345

processing. For example, a suprathreshold target is presented for 100 ms and is immediately followed by a patterned mask. The target is visible to the subject, but the mask is held to limit its availability to 100 ms (Coltheart, 1983; Enns & Di Lollo, 2000). If the mask were not presented, it is argued, the target would remain available to the subject in a short-term visual store (Sperling, 1963). But those who use a poststimulus mask in this way are entertaining an implicit model in which the mask is either superposed on the target in the store or else displaces it from the store. How secure are we in assuming that a poststimulus mask will terminate the short-term visual store, or "icon"?

To address this issue, we combine the technique of backward pattern masking with another technique, that of part report, which was used by Sperling to demonstrate the very existence of short-term visual storage (Sperling, 1960). In Sperling's experiments, an array of 12 letters is briefly presented and then removed and replaced by a blank screen. If the subject is asked to report as many letters as he can from the entire array, typical performance is in the range of 3–5 letters, accurately reported in the correct locations. If instead, a few hundreds of milliseconds after the disappearance of the letter display, a cue indicates which row of letters to report, then the subject can recover 3–4 letters irrespective of which row is specified. This result suggests that, at the time of presentation of the cue, a large proportion of the array is still available to the subject in iconic memory. The row to report is chosen at random, and part-report performance is taken to be the number of letters reported per row, multiplied by the number of rows. The advantage of part report over whole report provides a measure of *iconic storage*, and this is the operational definition that we use in the present paper.

A key question is whether a part-report advantage survives when a random-pattern mask is introduced between the target array and the cue. In the traditional view (Becker, Pashler, & Anstis, 2000; Sperling, 1960), the iconic storage of the target is completely terminated by the mask. Thus Gegenfurtner and Sperling (1993, p. 865), in an analysis of transfer from iconic to durable storage, write:

> Immediately after the cue, attention shifts to the cued row of the display . . . . From this moment on, until the poststimulus mask ends all iconic transfer, selective transfer occurs from the cued row.

If a cue follows the mask, it is an empirical question whether the cue will allow the subject to select from within the target array that was presented before the mask. An advantage of part report over whole report under these circumstances would provide operational evidence for a store that preserves a visual representation of the target but is resilient to the aftercoming mask.

Consider the predictions of the traditional models of iconic memory in such masking experiments (see Figure 1). The target is presented and enters the icon. Letters are then read out of the icon into a more durable store. In whole report, or in part report before the cue is presented, the read-out is nonselective and can be from anywhere in the array of target letters. When the mask is presented, letters in the durable store are safe, but everything from the icon is lost. So, when the cue is presented, attention is directed to the relevant position in the icon, but there are no surviving letters. Performance in partial report is therefore predicted to be proportionately the same as that in whole report, since it can be based only on information that is obtained during the nonselective read-out phase.

Traditional models of iconic storage assume that the representation of the target can be lost via "replacement" or via "integration" with the mask. In the "replacement" model, iconic memory is a screen that is cleared by the arrival of each fresh stimulus. In the "integration" model, iconic memory is a screen on which successive items are superposed, or integrated. Whatever the exact mechanism of loss, we wish to question the implicit acceptance of such models by experimenters who use a pattern mask to terminate perceptual processing.

We have measured part-report performance with a cue that followed a pattern mask. Clearly, we could unfairly ensure survival of the iconic

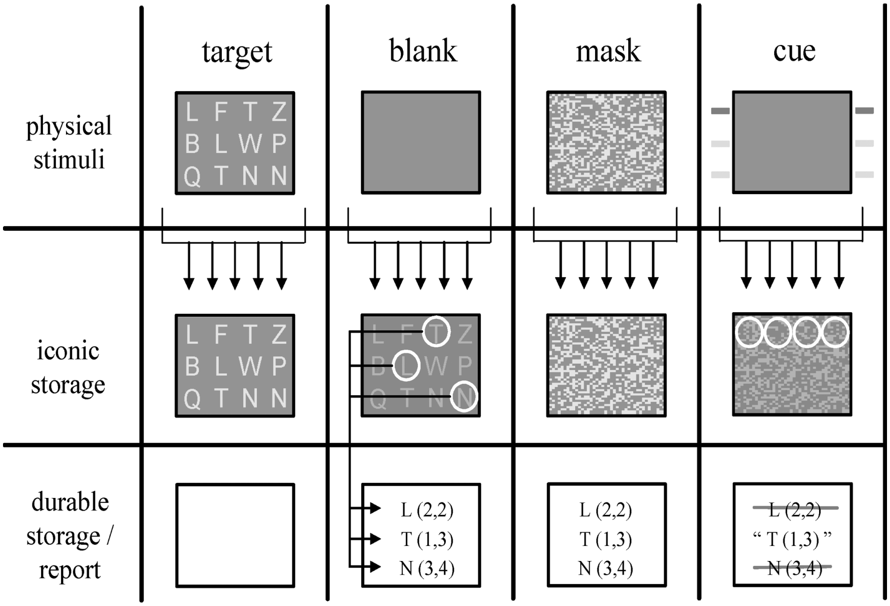

Figure 1. *A conventional view of the role of iconic storage in a masked part-report trial. Successive events are represented from left to right, and different levels of representation are shown in different rows. In the first column, the target is presented and enters the icon. In the second column, letters are read out of the icon into a more durable store. In whole report, or in part report before the cue is presented, the read-out is nonselective and can be from anywhere in the array of target letters (here, selected letters are circled). In the third column, the mask is presented. Letters in the durable store are safe, but everything from the icon is lost. In the fourth column, the cue is presented. Attention is directed to the relevant position in the icon, but there are no surviving letters. Performance in partial report is therefore predicted to be proportionately the same as that in whole report, since it can be based only on information that is obtained during the nonselective read-out phase (second column).*

representation of the target by using a masking stimulus so weak that its presentation was trivial. To avoid this we have chosen a pattern mask whose spatial frequency, contrast, and intensity were such that, when superposed on the target array, it rendered the target letters completely illegible. Rather than performing the superposition by calculation we have exploited the limited temporal resolution of the early visual system. For each subject, in trials that were counterbalanced with the main experiment, we checked that identification of the target letters was reduced to chance when target and mask were presented on consecutive screen-refreshes of a CRT. If a pattern mask with these properties follows the target after a delay, and if a subsequent cue still gives a part-report advantage, then we must question the use of masking to "terminate the icon".

We also systematically varied the interval between the target and the cue, while keeping constant the interval between target and mask. There are two reasons for such measurements. First, it is important to show that any difference between part report and whole report does decline with time: Otherwise it could be argued that the part-report advantage arises from competition at a response stage, since there are more items to report in the whole-report condition (von Wright, 1972). Second, by varying the cue delay, we obtain estimates of the duration of the hypothetical store. These are lower bound estimates, owing to the time required to process and interpret the cue (Averbach & Coriell, 1961; Coltheart, 1983).

Methods

Observers

Three observers participated in these experiments, the two authors (HES and JDM) and a third observer (RES) who was naïve to the purposes of the present experiment. All observers were highly practised. The part-report procedure requires the observer to follow a relatively complicated strategy of using a cue to improve performance. Unfortunately, an observer can choose not to use a cue that he deems to be irrelevant. To obtain reliable part-report effects it is therefore important to have explicit instructions and extensive practice (Averbach & Coriell, 1961; Chow, 1986; Sperling, 1960). Moreover, it is also known that performance under conditions of backward masking improves over many trials (Karni & Sagi, 1993; Wolford, Marchak, & Hughes, 1988). Each observer served in 12 practice sessions (i.e., 2,160 part-report trials), plus a further 8 or 12 experimental sessions.

Apparatus

Stimuli were presented on a 17-in. Sony FD Trinitron monitor driven at a frame rate of 100 Hz by a Cambridge Research Systems (CRS) VSG 2/3 graphics card housed in a Pentium computer. The display was gamma corrected (linearized) from measurements made with a CRS ColorCAL colorimeter. Stimulus presentation and response recording were controlled via custom software written in the C++ programming language. Observers' responses were entered via the computer keyboard. Observers performed the task in a dimly lit laboratory at a viewing distance of 1.5 m.

Stimuli

Figure 2 provides a schematic representation of our stimuli. The display was set to a uniform grey background (chromaticity CIE, $x = 0.33$, $y = 0.33$, luminance 19.5 cd/m^2). Target, mask, and cue contrasts were positive and matched at 58%. Targets were presented within one 10-ms video frame (at a refresh rate of 100 Hz); masks and cues were longer and spanned two frames.

Target stimuli were three-by-four arrays of letters rendered in Arial font. The height of each letter was fixed at 50 pixels (or 8.4 mm on our display) and subtended 0.64 degrees of visual angle at the viewing distance of 1.5 m. Each letter was centred in a square area of 0.77 by 0.77 degrees of visual angle, and 12 such squares (three rows and four columns) comprised the stimulus area.

On each trial, 12 target letters were chosen at random from 15 possible letters (B C D F H L

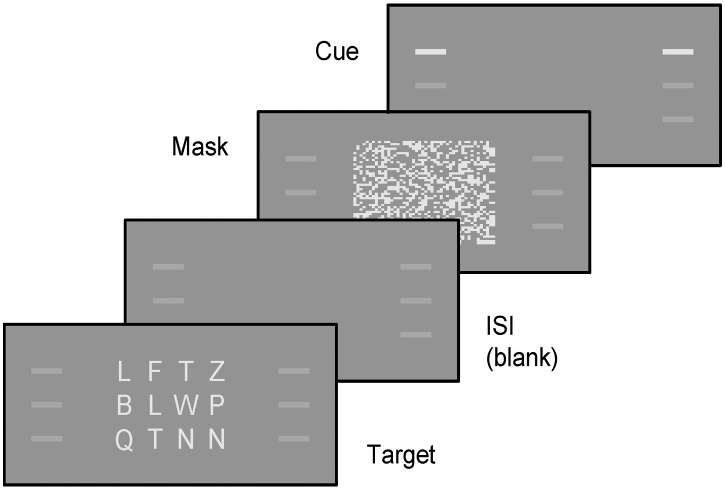

Figure 2. *Schematic representation of the stimuli in a part-report trial. A target array of 12 letters is briefly presented (<10 ms), followed after long (100-ms) or short (0-ms) ISI by a checkerboard pattern mask (<20 ms), and then followed by a part-report cue (<20 ms). In this example the cue indicates that the top row should be selected for report. Skeleton cue markers are present throughout the trial.*

N P Q S T V W X Z). In order to minimize the likelihood of generating pronounceable letter strings we chose to exclude vowels and the letter “Y” from our arrays (Sperling, 1960). On the basis of data from preliminary experiments, we also excluded 5 additional letters (G J K M R) to eliminate the most visually confusable pairs. On each trial, individual letters were chosen with replacement, so chance performance corresponded to 1/15 letters per letter guessed.

Mask stimuli were random 2-bit checkerboards with 50% of checks set to the background level and the remainder incremented to obtain a contrast of 58%. Individual checks were 4 pixels high and 6 pixels wide to match the horizontal and vertical stroke-widths of the letters. Checkerboards were 45 by 45 checks and subtended 2.3 by 3.5 degrees of visual angle. The mask area was thus larger than the target area and overlapped the target area on all sides.

Cues, to indicate the row to be reported in the part-report conditions, were rectangular markers (0.77 by 0.13 degrees of visual angle), presented to the left and right of the target area. They were aligned horizontally with the middle of the relevant row and were separated from the target area by 1.2 degrees of visual angle. Cues were set to the same positive contrast as that for the target and mask. The cued row was chosen at random, with equal likelihood for each of the three rows and an explicit constraint that each row was cued the same number of times in an experimental session. Throughout the experiment, faint rectangular bars (20% contrast) were displayed in the possible cue locations to provide a skeleton frame of reference.

Procedure

On each part-report trial, we presented a target array of letters, followed by a checkerboard mask and later followed by a cue. For whole report, there was no cue. There were nine stimulus conditions:

1. Whole report, target-to-mask interstimulus interval (ISI) = 0 ms, no cue.
2. Whole report, target-to-mask ISI = 100 ms, no cue.
3. Part report, target-to-mask ISI = 0 ms, mask-to-cue delay = 0 ms.
4. Part report, target-to-mask ISI = 100 ms, mask-to-cue delay = 0 ms.
5. Part report, target-to-mask ISI = 100 ms, mask-to-cue delay = 100 ms.
6. Part report, target-to-mask ISI = 100 ms, mask-to-cue delay = 200 ms.
7. Part report, target-to-mask ISI = 100 ms, mask-to-cue delay = 300 ms.
8. Part report, target-to-mask ISI = 100 ms, mask-to-cue delay = 400 ms.
9. Part report, target-to-mask ISI = 100 ms, mask-to-cue delay = 500 ms.

For the short-ISI conditions (ISI = 0 ms), target and mask were presented on consecutive frames. For the long-ISI conditions (ISI = 100 ms), there were 10 blank (uniform grey background) frames between target and mask presentations. The six cue delays used in the long-ISI conditions corresponded to 120, 220, 320, 420, 520, and 620 ms after the target offset. Data for each condition were obtained in unmixed blocks. Blocks for different conditions were interleaved and counterbalanced within and across observers.

Observers initiated each trial with a key press. Three 100-ms auditory tones were played, and the sequence of visual stimuli began. In whole-report conditions, observers were asked to report as many letters as possible from the target array. They were asked to initially record their responses with pencil and paper and then to type their responses into the computer. They were required to enter all 12 letters and were encouraged to guess if they were unsure. In part-report conditions, observers were asked only to report the four letters from the cued row. In all cases, the number of letters correct was defined as the number of letters correctly identified in the correct locations. After entering their response, observers were given feedback. The target array (whole report) or the target row (part report) was shown on the left of the display, and the observer’s response was shown on the right along with a summary specifying the number of letters correct. After seeing the feedback, observers could initiate the next trial with a key press. Each experimental session consisted of 60 trials for a particular condition. The nine conditions were repeated six times (observers RES and JDM) or four times (observer HES).

Results

Does the part-report advantage survive if the cue follows a pattern mask?

The data presented in Figure 3 were obtained in Conditions 1, 2, 3, and 4 (see Procedure). Values on the ordinate are estimates of the number of letters available to the observer from each presentation. In whole-report conditions (dark-grey bars) this is simply the average number of letters correctly reported. In part-report conditions (light-grey bars) we assume (Sperling, 1960) that the letters correctly reported from the cued row represent a random sample of the letters available to the observer at the time of the cue and that, for the entire array, the number of letters available is given by the number of letters reported per row, multiplied by the number of rows. The part-report data given here are for a cue delay of 0 ms after the offset of the mask.

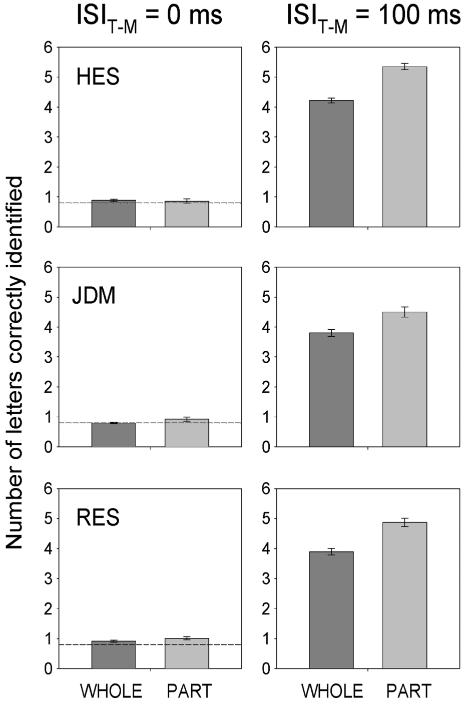

Figure 3. *A comparison of part and whole report at short and long ISI. Data are from Conditions 1, 2, 3, and 4 (see Procedure). Values on the ordinate are estimates of the number of letters available to the observer from each presentation. In whole-report conditions (dark-grey bars) this is the average number of letters correctly reported. In part-report conditions (light-grey bars) this is the number of letters reported per row, multiplied by the number of rows. Panels on the left show performance at target–mask ISI = 0 ms. The dashed lines represent chance performance (i.e., 12 × 1/15 = 0.8 letters for the whole array). Panels on the right show performance at target–mask ISI = 100 ms. Error bars show ±1* SE.

Panels on the left show performance at target–mask ISI = 0 ms. The dashed lines represent chance performance (i.e., 12 × 1/15 = 0.8 letters for the whole array). Multiple one-sample *t* tests confirm that none of these values is significantly different from chance. Here and in all subsequent analyses the critical alpha value is .05. After applying the Bonferroni correction, $p_{\text{crit}} = .05/6 = .008$ for each of the six *t* tests. Panels on the right show performance at target–mask ISI = 100 ms.

A 2 × 2 within-subjects analysis of variance (ANOVA) with ISI and report type as factors revealed a significant main effect of ISI, $F(1, 2) = 347$, $MSE = 37.5$, a significant main effect of report-type, $F(1, 2) = 157$, $MSE = 0.760$, and a significant interaction, $F(1, 2) = 23.2$, $MSE = 0.554$. Simple main effects of report type at short and at long ISIs were analysed by one-way ANOVAs and confirmed no significant part-report advantage at short ISI, $F(1, 2) = 1.91$, $MSE = 0.008$, but a significant part-report advantage at long ISI, $F(1, 2) = 53.3$, $MSE = 1.31$. Thus, when target and mask are presented on consecutive frames and integrated in the early visual system, we observe no part-report advantage, but when the same target and mask are separated by 100 ms, there is a part-report advantage.

Under whole-report conditions, observers predominantly reported letters from the middle row (on average 2.8 ± 0.2 letters compared to 0.6 ± 0.1 letters for top and for bottom rows). In the part-report conditions, the gain was primarily on the top and bottom rows, +0.6 ± 0.1 and +0.5 ± 0.1 letters, respectively, compared to −0.1 ± 0.2 letters for the middle row.

Does the part-report advantage decline if the cue is delayed?

The data presented in Figure 4 were obtained in Conditions 2, 4, 5, 6, 7, 8, and 9 (see Procedure). Small rectangles indicate stimulus timing—black rectangles for the mask and grey for the cue. Dashed lines show whole-report performance. The part-report advantage (i.e., the height of the solid line above the dashed line) clearly declines as the cue is delayed and is lost when the cue is presented approximately 700 ms after the target.

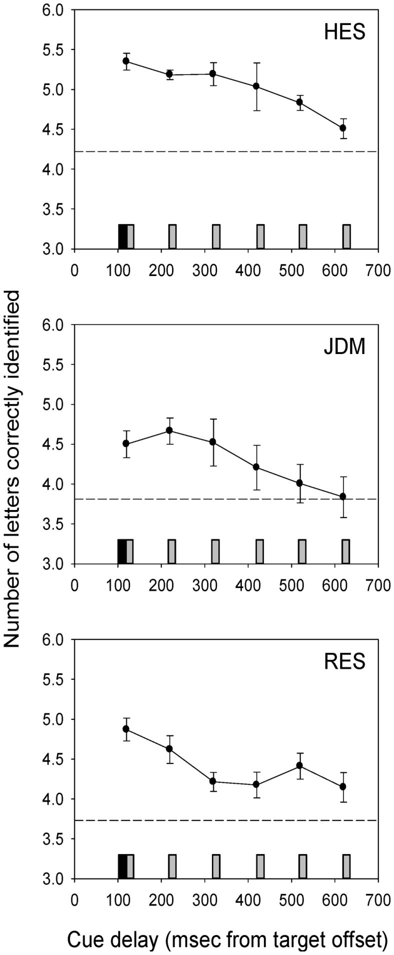

Figure 4. *Part-report advantage as a function of cue delay. Data are from Conditions 2, 4, 5, 6, 7, 8, and 9 (see Procedure). Small rectangles indicate stimulus timing—black rectangles for the mask and grey for the cue. Cue delay is expressed in ms after target offset. Dashed lines show whole-report performance. Data points show part-report performance. Error bars show ±1* SE.

A one-way, within-subjects ANOVA, with cue delay as a factor, confirmed a significant effect of cue delay, sphericity assumed, $F(5, 10) = 8.52$, $MSE = 0.229$, $p = .002$; Huynh–Feldt correction applied, $\varepsilon = .576$, $F(2.88, 5.76) = 0.852$, $MSE = 0.398$, $p = .02$. Linear trends analysis confirmed a significant linear component, $F(1, 2) = 154$, $MSE = 1.12$, to the decay. The part-report advantage declines as the time of presentation of the cue is delayed.

Are visual confusions more common in part report than in whole report?

An analysis of errors on a particular task can be instructive in determining how the information used to solve that task is encoded (Conrad, 1964). If the difference in performance between part and whole report were mediated via a short-lived store in which information is visually encoded, we should expect more "visual" confusions in part-report conditions.

In addition to recording the number of correct responses, we recorded response errors, and from these data we constructed two confusion matrices, one for part report (with the shortest cue delay) and one for whole report. We normalized the occurrence of each confusion pair (presented letter–reported letter) by the total number of presentations of the presented letter and the total number of reports of the reported letter. The first normalization has very little effect (since all letters were equally likely to be presented); the second normalization corrects for biases in guessing.

To quantify visual similarity, we defined an index based on the common presence, or common absence, of five visual features: vertical lines, horizontal lines, diagonal lines, curved contours, and closed forms.

As expected, confusion indices for part report (P) are correlated with confusion indices for whole report (W; $r_{PW} = .33$). Visual similarity indices (V) are correlated with confusion indices for part report ($r_{VP} = .42$) and with confusion indices for whole report ($r_{VW} = .25$). A statistical test for partial correlation (which allows us to account for the fact that r_{VP} and r_{VW} are not independent) shows that the correlation r_{VP} is significantly higher than the correlation r_{VW}, $t(207) = 2.3$. Thus, more similar letters were more likely to be confused, both in part and whole report, but the correlation between visual similarity and confusability was stronger in the part-report data than in the whole-report data.

Discussion

The nature of iconic storage

Our results are at odds with a view in which iconic storage is automatically displaced or overwritten by new information—the view in which the icon is a single screen on which successive images are projected. We have presented evidence for a neural representation of the target that decays within approximately 700 ms and is not completely destroyed by the delayed presentation of a mask. A spatial cue can apparently direct attention within this representation even when a pattern mask has come and gone.[1]

But what is likely to be the neural substrate of this representation? And in what format does it store visual information—as a bit-mapped array of pixels, as vectors representing edges and lines, as lists of abstract features, or as categorized objects? We can be sure that the representation is postreceptoral: The rods and cones themselves could not distinguish isomerizations produced by the target letters and the superposed mask (Burns & Lamb, 2004), and so it would be difficult to argue that the late part-report cue allowed recovery of information stored in the photoreceptors. Equally, however, we know that the representation cannot be postcategorical and nonvisual: The representation must retain information about visual position, since our part-report cue was a spatial one. Moreover, since errors in part report were more likely to be visual confusions than were errors in whole report, the representation must preserve elementary visual features.

How could the traditional account of the icon be modified to accommodate our results? In the following paragraphs, we set out two alternative hypotheses:

A. The selection in part report is from a store that the mask cannot penetrate. It is implausible to model the icon as a single store located at one site within the visual system. Rather, the internal representation of a visual stimulus is likely to be distributed in time at each of the several stages through which it passes: There is certainly persistence at the receptoral level, owing to the time constants of the rods and cones, but there is also a temporal distribution of later representations, where the stimulus is thought to be represented first as edges and textures and later as object descriptions that carry with them location tags. As the distributed representation of the target array moves through this sequence of analysis, it is pursued by a similarly distributed representation of the mask. But the analysers are also necessarily filters, and thus there will be a level of analysis to which the checkerboard pattern mask cannot pursue the targets—a level where letters are represented not as in terms of their local edges but in terms either of their high-level features or of their identities as particular characters. In so far as such representations preserve information about location, then a part-report cue may allow selection from these levels of storage. There is other experimental evidence for levels of storage that are increasingly tuned to a particular stimulus dimension (Lakha & Wright, 2004; Smithson, 2000) and for short-term visual storage at a level

---

[1] Gegenfurtner and Sperling (1993), although holding a traditional view of iconic storage, obtain a result similar to the present one: With cues that occur after mask onset, they find for one subject a weak advantage of part report compared to whole report, and for the other subject a larger part-report advantage (their Figure 9). Importantly, both of their subjects show a decline in part-report advantage with increasing cue delay, even when the cue follows the mask (their Figures 7 and 8).

where stimuli are represented as abstract features or indeed as objects that belong to particular categories (Di Lollo, 1980; Coltheart, 1983; Irwin & Yeomans, 1986; Mewhort, Campbell, Marchetti, & Campbell, 1981).

By this first hypothesis, then, an aftercoming pattern mask fails to "terminate the icon", because it cannot penetrate the higher levels of visual analysis, where target stimuli are represented in terms of abstract visual features or categorized objects and where nevertheless positional information is retained.

B. The selection in part report is from a four-dimensional store. The second hypothesis that we consider is a more radical one. In the case of hearing, it is assumed that stimuli that occur very closely in time can mask each other, but that once they are separated by 100 ms or more, consecutive items are held in a short-term auditory store that preserves their physical features and their sequence (Darwin, Turvey, & Crowder, 1972). A tape-recording provides the traditional metaphor for this immediate auditory memory. Why should not the model for visual short-term storage be a video tape-recording, with a sequence of frames? Can we really exclude the possibility that that the short-term visual store preserves successive events discretely in sequence in their sensory form, as the auditory store preserves a rapid sequence of phonemes? In this model, rather than being a two- or three-dimensional snapshot of the world, the icon is four-dimensional, representing time as well as the three dimensions of visual space. It can accommodate contents such as trajectories, optic flow of visual texture, complex gestures of the hands, and the lip movements that occur during speech. A part-report cue may allow attention to be selectively directed anywhere within the four dimensions of this perceptual store. The existence of such a store has been postulated previously (Allik & Bachmiann, 1983; Johansson, 1983; Klatzky, 1983; Mollon, 1969; Phillips, 1983), and an explicit and detailed model has been developed by Schill and Zetzsche (1995).

By this second hypothesis, an aftercoming mask fails to "terminate the icon", because the mask and the target are independently represented within a four-dimensional store.

Implications for studies that employ backward masking

Understanding the way in which the visual system handles rapidly sequential stimuli is important in understanding vision itself, but also has implications for a large body of work in psychology. Psychologists commonly employ masks for one of three reasons: (a) to generate subliminal stimuli, (b) to degrade the stimulus so that performance falls into a measurable range, or (c) to limit the amount of time for which a stimulus is available for perceptual processing.

The present study does not bear on the first of these applications, since at long ISI our target stimuli were above threshold and clearly visible. With regard to the second application, we do not dispute that backward masking can impair performance relative to performance with the target alone. By Hypothesis A (above), masking occurs either because some filtered elements of the mask do penetrate as far as the high-level store that contains the representation of the target or because the mask diverts resources from analysis of the target. By Hypothesis B, masking may be due either to limited temporal resolution within the four-dimensional representation or to deflection of processing resources (Schill & Zetzsche, 1995). It is because the processes of backward masking are not well understood that its casual use should be avoided by psychologists, a point forcibly made by Eriksen (1980). One specific danger is that the chosen mask may not affect all stimuli equally and so will confound differences between experimental conditions.

Our present experiments directly address the third application of masking. The almost universal assumption is that presentation of a patterned mask will terminate the icon, and psychologists therefore use masks to limit the time during which target stimuli are available to the subject. If a delayed spatial cue can allow selective report from within a masked target array, then we must reevaluate many studies in which visual masking is used to terminate the perceptual processing of earlier items. We illustrate our argument with examples from two different areas of psychological research: reading and attention.

Van Orden (1987) used a masking experiment to argue that skilled readers use phonological features in visual identification of words: When target words were followed by a pattern mask, subjects continued to show false positives to phonologically similar foils but not to orthographically similar foils. By assuming that an aftercoming pattern mask terminates visual processing, Van Orden was led to the counterintuitive claim that the most *rapidly* activated codes in reading were phonological ones. Our own proposal would be that orthographical confusions are not observed under masked conditions because the mask has its effect by flooding the orthographic analysers with noise, not by limiting the time for which the target word is available.

Saarinen and Julesz (1991) used a masking experiment to measure the speed of attentional shifts across the visual field. A short sequence of numerals was presented one by one at random retinal locations, and observers were required to identify the sequence. Each numeral was followed by a pattern mask, and it was assumed that to identify the sequence correctly the observers had to move their attention at the same speed as the presentation rate of the stimuli. If masks do not terminate the icon, Saarinen and Julesz's estimate of the speed of attentional shifts must be revised.

CONCLUSIONS

The part-report procedure has been taken to give an operational measure of iconic storage, and pattern masks have been taken to terminate this iconic storage. While it is beyond doubt that there is more than one visual memory, pattern masks are still widely used by experimental psychologists to limit (to, say, 100 ms) the time for which suprathreshold targets are "available for processing". The purpose of our paper is to prompt discussion of what it means to say a mask terminates perceptual processing and what might be the nature of the store that allows the continuing part-report advantage. One thing is eminently clear: It is not possible concurrently to define the duration of visual short-term storage by part report and to hold that pattern masks terminate this storage.

Original manuscript received 12 January 2005
Accepted revision received 28 June 2005

REFERENCES

Allik, J., & Bachmiann, T. (1983). How bad is the icon? *Behavioral and Brain Sciences, 6*, 12–13.

Averbach, E., & Coriell, A. S. (1961). Short-term memory in vision. *The Bell System Technical Journal, 40*(1), 309–328.

Becker, M. W., Pashler, H., & Anstis, S. M. (2000). The role of iconic memory in change-detection tasks. *Perception, 29*(3), 273–286.

Burns, M. E., & Lamb, T. D. (2004). Visual transduction by rod and cone photoreceptors. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences.* Cambridge, MA: MIT Press.

Chow, S. (1986). Iconic memory, location information, and partial report. *Journal of Experimental Psychology: Human Perception and Performance, 12*(4), 455–465.

Coltheart, M. (1983). Iconic memory. *Philosophical Transactions of the Royal Society, B, 302*(1110), 283–294.

Conrad, R. (1964). Acoustic confusions in immediate memory. *British Journal of Psychology, 55,* 75–84.

Darwin, C. J., Turvey, M. T., & Crowder, R. G. (1972). An auditory analogue of the Sperling partial report procedure: Evidence for brief auditory storage. *Cognitive Psychology, 3*, 255–267.

Di Lollo, V. (1980). Temporal integration in visual memory. *Journal of Experimental Psychology: General, 109*(1), 75–97.

Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences, 4*(9), 345–352.

Eriksen, C. W. (1980). The use of a visual mask may seriously confound your experiment. *Perception and Psychophysics, 28*(1), 89–92.

Eriksen, C. W., & Collins, J. F. (1967). Some temporal characteristics of visual pattern perception. *Journal of Experimental Psychology, 74*(4), 476–484.

Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance, 19*(4), 845–866.

Irwin, D. E., & Yeomans, J. M. (1986). Sensory registration and informational persistence. *Journal of Experimental Psychology: Human Perception and Performance, 12*(3), 343–360.

Johansson, G. (1983). Optic flow, icons, and memory. *Behavioral and Brain Sciences, 6*, 23–24.

Karni, A., & Sagi, D. (1993). The time course of learning a visual skill. *Nature, 365*(6443), 250–252.

Klatzky, R. L. (1983). The icon is dead: Long live the icon. *Behavioral and Brain Sciences, 6*, 27–28.

Lakha, L., & Wright, M. J. (2004). Capacity limitations of visual memory in two-interval comparison of Gabor arrays. *Vision Research, 44*(14), 1707–1716.

Mewhort, D. J., Campbell, A. J., Marchetti, F. M., & Campbell, J. I. (1981). Identification, localization, and "iconic memory": An evaluation of the bar-probe task. *Memory and Cognition, 9*(1), 50–67.

Mollon, J. D. (1969). *Temporal factors in perception.* Unpublished D. Phil. dissertation, University of Oxford, Oxford, UK.

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

Palmer, J. (1988). Very short-term visual memory for size and shape. *Perception & Psychophysics, 43*, 278–286.

Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics, 16*(2), 283–290.

Phillips, W. A. (1983). Change perception needs sensory storage. *Behavioral and Brain Sciences, 6*, 35–36.

Phillips, W. A., & Baddeley, A. D. (1971). Reaction time and short-term visual memory. *Psychonomic Science, 22*, 73–74.

Saarinen, J., & Julesz, B. (1991). The speed of attentional shifts in the visual field. *Proceedings of the National Academy of Sciences (USA), 88*(5), 1812–1814.

Schill, K., & Zetzsche, C. (1995). A model of visual spatio-temporal memory: The icon revisited. *Psychological Research, 57*(2), 88–102.

Smithson, H. E. (2000). *Visual masking.* Unpublished PhD dissertation, University of Cambridge, Cambridge, UK.

Sperling, G. (1960). The information available in a brief visual presentation. *Psychological Monographs: General and Applied, 74*(11), 1–29.

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

Treisman, A., Russell, R., & Green, J. (1975). Brief visual storage of shape and movement. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and performance V* (pp. 699–721). London: Academic Press.

Van Orden, G. C. (1987). A ROWS is a ROSE: Spelling, sound, and reading. *Memory and Cognition, 15*(3), 181–198.

von Wright, J. M. (1972). On the problem of selection in iconic memory. *Scandinavian Journal of Psychology, 13*(3), 159–171.

Wolford, G., Marchak, F., & Hughes, H. (1988). Practice effects in backward masking. *Journal of Experimental Psychology: Human Perception and Performance, 14*(1), 101–112.