PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Iconic Memory [and Discussion]

M. Coltheart, D. R. J. Laming, D. A. Routh and D. E. Broadbent

Phil. Trans. R. Soc. Lond. B 1983 **302**, 283-294
doi: 10.1098/rstb.1983.0055

'. Trans. R. Soc. Lond. B **302**, 283–294 (1983)
ıted in Great Britain

Iconic memory

By M. Coltheart

Department of Psychology, Birkbeck College, Malet Street, London WC1E 7HX, U.K.

Investigations of the processing of brief visual displays, and the explanation of such processing in terms of iconic memory, are reviewed. It is concluded that the concept of a pre-categorical sensory memory for visual material remains tenable. The ability to report material from brief visual displays is seen as depending upon parallel (and perhaps unlimited) transfer from iconic memory to a post-categorical memory mode, followed by a limited (and perhaps serial) transfer to an output stage. Decisions about, or responses to, items can only be made when they are in the output stage. Because transfer out of the post-categorical mode can be performed on the basis of pre-categorical stimulus features, pre-categorical information about items in the post-categorical mode must be accessible. One way in which this would be possible is if the transfer of an item into the post-categorical mode takes the form of the creation, to represent the item, of a temporary file of information including both pre-categorical and post-categorical features of the item. Any such feature can be used as the basis for selecting the item for transfer from the post-categorical mode to the output stage, for subsequent decision or report.

The human observer's ability to process briefly presented alphanumeric displays has both impressive and unimpressive aspects. What is impressive is the level of performance attainable even when such displays are extremely brief: we can read a word when it is illuminated by a bolt of lightning or by a stroboscope; under such conditions of illumination the duration of the display is extremely small: less than a millisecond. What is unimpressive, on the other hand, is what happens when the brief visual display consists of a number of separate stimuli – a matrix of a dozen unrelated letters, for example – all of which are to be reported if possible. Here rather few items can be reported: about four or five items, on the average.

These two basic properties of our ability to process visual displays were first established in the nineteenth century, and have been repeatedly verified during the past 20 years. Any theory of the processing of brief visual displays must therefore offer explanations of both observations if it is to have any pretensions to adequacy.

It would seem, from the fact that a word or a letter can be identified even when displayed for less than a millisecond, that some form of buffer storage for brief visual displays must exist. No one would be willing to claim that all of the processes needed for identifying a word or a letter could be carried out in so short a time. It follows that there must be a period of time after the offset of very brief displays during which (*a*) the stimulus has not yet been identified, and yet (*b*) the stimulus is no longer physically present. During this period, then, the stimulus, even though no longer physically present, must still be available as a source of input to the machinery of letter and word identification. It must, in other words, be represented in a form of memory holding pre-categorical (not yet identified) information. It is evident, from the fact that when many items are presented only few can be reported, that the system of processes by which we deal with brief visual displays must at some stage possess a bottleneck imposing limitations upon the number of items with which the system can deal effectively.

Thus we can be a little more specific concerning what needs to be accomplished if we are to formulate an adequate theory of the processing of brief visual diaplays. We will, at the very least, need to offer an account of the nature of the pre-categorical buffer memory that must constitute an early stage in such processing, and also an account of the nature of the bottleneck that must exist at some stage in this processing.

Let us begin by considering possible explanations for the inability of observers to report all the items from a display that is brief and contains many items. Appropriate data are those of Sperling (1960, fig. 4). Four subjects were used. Displays lasted for 15 ms and consisted of six letters (two rows of three). Subjects were asked to report as many letters from the display as they could. The mean number of correctly reported letters was almost exactly 4.0 for three of the subjects and was slightly less than 5.0 for the remaining one, letters being scored as correct only if they were written in the correct locations in 3×2 response grids. Thus, even when there are as few as six letters in the display, not all can be reported correctly. Why?

The most immediately obvious explanation for this limitation of performance is that the display was so brief. It might be, for example, that character identification is a serial process and 15 ms is insufficient time for the performance of six identifications. We can rule this explanation out immediately, however, because Sperling (1960, fig. 4) showed that giving subjects much more time (increasing display duration from 15 ms to 500 ms) did not improve performance at all: even with these very long displays, subjects still averaged fewer than five letters correctly reported from the six-letter displays.

Perhaps instead the explanation is that there exists a limitation upon the number of items that can be *simultaneously perceived*. Perhaps the perceptual apparatus, though operating in parallel across items rather than serially, can only deal with four or five items at a time and having done so is subsequently in a refractory state for a period. The partial report technique introduced by Sperling yielded results permitting us to dismiss this explanation. Using displays of twelve letters (three rows of four), Sperling found that if the subject was asked to write down the letters from only one row, the row required being specified by an auditory cue presented immediately after display offset, the mean number of letters correctly reported in their correct positions was 2.7/4 for the worst of five subjects and 3.67/4 for the best, the mean across subjects being 3.03/4. Because the subjects could not know until after display offset which row was to be reported, they must have had available, in some form of storage, at least 3.03 letters from each row, on the average; that is on average at least 9.1 letters from the twelve-letter display must have been represented in some form of post-display storage. Under the same stimulus conditions, when these subjects were asked to make a full report, they averaged only 4.3 letters correct.

This limitation in the number of display items reported cannot have been due to a limitation in the number of items that can be simultaneously perceived, because the results obtained with the partial report technique show that at least 9.1 letters, on the average, are perceived well enough for them to be represented after display offset in a form of storage that can be used to perform the partial report task.

There are two reasons for believing that this value, 9.1/12, is, in fact, an underestimate of the total number of items perceived from the display and stored after its offset. Firstly, because responses were scored as correct only if the correct position as well as the correct item was given, the number of display items stored with *or* without information about items' locations in the display may have been greater than 9.1. Secondly, because it can take as long as 200 ms or more for the subjects to process the partial report cue to the point where it is known which row is to be reported (see, for example, Averbach & Coriell 1961), the subjects in Sperling's partial report experiment will have been retrieving information from the post-display storage after it has been in existence for more than 200 ms; and if this form of storage exhibits rapid decay (as it appears to: this is discussed below), then the value 9.1/12 represents an estimate of the capacity of this storage mode after it has partly decayed, and hence is an underestimate of the true capacity.

Let us now consider a third possible explanation of the limitations on full report from tachistoscopic displays, namely, *output interference*. Perhaps nearly all, or even all, the items from such displays are perceived and stored in memory, but it is only possible to report a subset of the memory items, because the act of reporting has a destructive effect upon memory (by, for example, preventing rehearsal of the items not yet reported). Perhaps as many as twelve or more items can be held in memory until the subject begins to report them; then, by the time four or five are reported, the remainder have become unreportable through the action of output interference.

This third explanation is refuted by the results obtained by Sperling in an experiment in which the post-display cue was presented at various times after display offset, rather than always immediately after it. With twelve-letter displays (three rows of four letters) and a cue immediately after display offset, subjects averaged about three letters correct out of four. As the cue was delayed, however, this value steadily decreased until, with a cue presented 1 s after display offset, performance averaged only 1.5 correct out of four. Multiplying these values by three (since there were three rows presented), we may infer that shortly after the display at least 9 letters were available in memory, whereas 1 s or so after the display only 4.5 letters were available. If the principal limitation on the number of items that can be reported is output interference, there is no reason why the number of items correctly reported in the partial report condition should be so small at the long cue delay. There is no reason why output interference should exert stronger effects when report is delayed; indeed, we know it does not, because when subjects are asked to report all the items they can (full report rather than partial report) they perform just as well when report is delayed as when it is immediate (Sperling 1960).

We have seen that three plausible explanations of the limitations evident in studies of tachistoscopic report may be dismissed on the basis of various results reported by Sperling (1960). A much more adequate explanation of all these results is offered by a theory (see, for example, Sperling 1967; Coltheart 1972, 1977) incorporating two distinct forms of memory: an initial high-capacity sensory buffer and a subsequent limited-capacity non-sensory memory system.

This explanation runs as follows. The visual display is initially represented in a high-capacity, rapidly decaying, specifically visual form of memory, usually referred to as iconic memory. The iconic memory of a display is created during the first few milliseconds or tens of milliseconds of the display (this time may depend on display energy, with the time needed to create iconic memory fully varying inversely with intensity). Once iconic memory is fully established, the presence of the display is irrelevant: performance will not be affected by switching off the display at this point because all subsequent processing is of the contents of iconic memory not of the display itself. This is why performance is independent of display duration, provided that display energy is sufficient for the creation of iconic memory.

This interpretation of the independence of tachistoscopic report and display duration might

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

be challenged by arguing that perhaps 15 ms is enough time to identify four letters, and a bottleneck higher in the system will not accept more than four letters, so that adding an extra 485 ms to the display cannot improve performance (that is, performance will be independent of display duration in the range 15–500 ms). This alternative account is, however, difficult to reconcile with certain effects of backward visual masking. If it is true that 15 ms after display onset four letters have already been identified (that is, translated into a non-visual code and so not susceptible to influence from a backward visual mask) then presenting a bright flash of light or other such mask at this point should not reduce the number of letters reported. This prediction is incorrect. A backward mask presented 15 ms after display onset very severely reduces the number of letters that the subject can report (Sperling 1963). Therefore, at this point after display onset some letters are:

(*a*) stored in a memory mode that is susceptible to visual masking, and hence not a post-identification stage, and

(*b*) will be identified and will contribute to eventual report if not masked.

Iconic memory, then, is established very early on in the lifetime of the visual display; after it is established the display no longer has any contribution to make to performance and can be removed without detrimental effect unless it is replaced by a masking stimulus, which can degrade or eliminate iconic memory.

For a number of reasons, tachistoscopic reports cannot be made solely on the basis of iconic memory. One reason is that the duration of iconic memory is shorter than the latency of written or spoken responses. A second reason is that items are represented in a pre-categorical form in iconic memory, and if an item is to be reported it must be represented in a post-categorical form (a phonological code if responding is spoken, a graphomotor code if responding is written). Hence the postulation of a second memory component, with negligible or zero decay, is needed if items are to last in some form of memory for more than a few hundred milliseconds after the display. I shall refer to this second component as 'durable storage' (Coltheart 1977). If an item in iconic memory is to be reported, it must first be transferred to durable storage.

The assumption is made that the transfer to this post-categorical stage, durable storage, is severely limited in the number of items with which it can deal, or else the capacity of durable storage is low. Thus even if iconic memory holds many items, only a few of them can attain durable storage and hence be reported. That is why only four or five items are produced when subjects engage in the full report task.

A further assumption is that the post-display cue in partial report experiments is applied to the contents of iconic memory, and so can control *which* of the many items in iconic memory are transferred to durable storage. That is why performance in partial report experiments is good: the transfer process, or the capacity of durable storage, is not overtaxed, because transfer is not attempted for many of the items in iconic memory.

The transfer process can operate under the control of any kind of cue provided that it refers to some pre-categorical form of information (spatial location, colour, size, brightness, orientation, etc.), because iconic memory contains only pre-categorical information.

Finally, it is assumed that iconic memory decays rapidly, with a lifetime of a few hundred milliseconds. This explains why performance in partial-report experiments declines as cue delay increases: the longer the interval is between display and cue, the more likely that some or all of the cued items will have decayed away from iconic memory before the cue is applied to iconic memory.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

This, then, is what one might call the standard theoretical account of the processing of brief alphanumeric displays. It offers explanations of a variety of basic findings, particularly:

(*a*) the limitations of full report;

(*b*) the superiority of partial report over full report;

(*c*) the decline in partial report performance as a function of cue delay;

(*d*) the lack of effect of exposure duration on tachistoscopic report;

(*e*) the effects of a backward visual mask.

The first to promulgate a theory of this form was Sperling (1963, 1967); other accounts of this kind of theory are given by Coltheart (1972, 1977); and there has been widespread dissemination of the theory in introductory textbooks over the past decade.

At a previous meeting of the Royal Society in 1979, on the psychology of vision, I mentioned, very briefly and very tentatively, certain issues that suggested to me that this standard account of the processing of brief visual displays required certain revisions. Results obtained by Eriksen & Eriksen (1974), van der Heijden (1978) and Allport (1977) could be taken as indicating (*a*) that access to a post-categorical memory system may *not* be difficult to achieve and may *not* be subject to severe capacity limitations, and (*b*) that the selection process operating in partial report experiments may be a selection not from items in pre-categorical storage but from items in post-categorical storage.

Hence 'a possibility is that iconic memory, rather than representing an early stage in the information processing sequence (a peripheral sensory buffer store), arises at a very late stage, after stimulus identification' (Coltheart 1980, p. 67). Since that was written, much has been learned about how the standard account needs to be revised and also about how it can be made more precise.

This recent work does not challenge the basic distinction between a pre-categorical sensory-memory system and a subsequent post-categorical system, but it does raise a number of questions about the ways in which these two systems have been conceived. I shall be considering some of these questions.

Can responses be made directly from the post-categorical stage?

In the theory originally proposed by Sperling (1967), there were, in fact, three components, rather than simply two. These were iconic memory, a post-categorical 'recognition buffer memory' holding post-categorical item representations, and an 'auditory information store'. Sperling's reason for adopting this three-component model was that he had deduced from experimental results that items were transferred out of iconic memory at a rate of about three items per 50 ms. This raised a problem, because Sperling believed that spoken responses were generated from an auditory information store the input to which was subvocal rehearsal, and subvocal rehearsal proceeds at a rate much slower than three items per 50 ms. Therefore he interposed between iconic memory and the output store a buffer store, and proposed that transfer of an item from iconic memory consisted of retrieval from long-term memory of a program of motor instructions for pronouncing that item. These programs were stored in the recognition buffer memory. The execution of these programs constituted transfer to the third component (the output store).

The division of the post-categorical stage into a recognition buffer and an output store has generally been neglected in subsequent theoretical treatments, but a very similar distinction has recently come to the fore, for example in the work of Duncan (1980, 1982). Duncan used the term 'first-level representation' to refer to a stage holding post-categorical item representations. He proposed that the attainment of this stage, even though it constitutes identification of an item, cannot be used directly for making responses to the item. If a response is to be made to the item (e.g. if it is to be reported, or if any decision is to be made about it), it must first be transferred to a 'second-level representation' and this transfer is accomplished by a limited-capacity, possibly even a serial, mechanism. Analysis of the contents of this second-level representation is what leads directly to responding.

In support of this view, Duncan (1980) reported results of visual search experiments. In one, for example, subjects were presented with four simultaneous alphanumeric stimuli, at the ends of the vertical and horizontal arms of an imaginary cross. The task was to press one button if a digit occurred in the vertical arm, to press a second button if a digit occurred in the horizontal arm, and to press both buttons if both arms contained a digit. Duncan found that detection of a digit in the vertical arm was better when the response to the horizontal arm was a correct rejection than when it was a hit, and detection of a digit in the horizontal arm was better when the response to the vertical arm was a correct rejection than when it was a hit. He interpreted these results as reflecting limitations on transfer from a post-categorical stage ('first-level representation') to an output stage ('second-level representation') in which decisions as to the presence or absence of a digit were made. All four items are assumed to attain the initial post-categorical stage. Then the digits represented in this stage (if there are any) are transferred to the output stage. Because of the limitations on the capacity of the transfer process, transfer is more difficult when there are two digits than when there is one. Therefore detection of a digit in one arm is more difficult when the other arm contains a digit too.

It is hard (though perhaps not impossible) to explain these results if one argues that responses can be made directly from the post-categorical stage, because one cannot claim that a digit is less likely to reach this stage when the display contains another digit: this claim would amount to assuming that the category to which a stimulus belongs can have effects before the stimulus attains categorization.

Further evidence is provided by Duncan (1982) from an experiment in which six-item displays consisting either of one digit and five letters, or else of three digits and three letters, were used. In a partial report condition, the task was to report all the digits present. In a full report condition, the task was to report all the items in the display. The mean number of digits reported was higher in the partial report than in the full report condition, and this was true even when masking was used to reduce the number of items reported in the full report condition to between two and three, making output interference an unlikely explanation of the results. Duncan also found that a digit was more likely to be reported in the partial report condition when it was the only digit in the display than when there were two other digits also present. These findings were interpreted by Duncan as indicating that 'in partial report, the main competition for the limited-capacity system comes from characters in the target class, non-targets being rejected at a prior level'. Because the distinction between target and non-target could only be made after categorization, this conclusion means that capacity limitations exert effects after categorization, and therefore that responses are not made directly from the initial post-categorical stage but from a subsequent output stage.

Where does selection operate in partial-report experiments?

If one adopts a *three*-component theory (pre-categorical sensory memory to post-categorical stage to output stage), there are now two possible loci at which partial report cues might be applied. The traditional view is that the partial report cue controls the transfer from sensory memory to the post-categorical stage. An alternative that can now be considered is that the partial report cue acts instead upon transfer from the post-categorical stage to the output stage.

Some early, and rather neglected, observations provide a little evidence in favour of the second alternative. Eriksen & Rohrbaugh (1970) and Dick (1974) reported that, in partial-report experiments where a post-display visual cue specifies one item for report from the display, subjects who fail to report the cued letter give very often as their response a letter that had been immediately adjacent to the cued letter in the display. It was shown that this error does not arise because of failures to appreciate which spatial location the cue indicated. If the cue controls transfer from iconic memory to a post-categorical stage, one might expect

(*a*) that the cued item has not decayed from iconic memory when the cue was applied to it (in which case the response should be correct), or

(*b*) that the item has partly decayed and is therefore difficult to identify (in which case errors should be identification errors based on partial visual information, and there is no reason why the response should tend to be another item from the display, especially not a closely adjacent item), or else

(*c*) that the item has completely decayed (in which case, again, there is no reason why the response made should be an adjacent item from the display).

The result suggests that there is a partial independence between information about what items had been in the display and where they had been located in the display (Cumming & Coltheart 1969), so that it is sometimes possible to have exact information about item identity with only approximate information about item location. This independence of item and location information is difficult to reconcile with the idea that the location cue is applied to a decaying image-like sensory trace akin to a progressively fading photograph. Further discussion of this point may be found in Coltheart (1983).

Additional evidence of this independence was provided by Townsend (1973). She varied the delay between display offset and the visual cue specifying a spatial location from which a letter had to be reported, and found, as did Averbach & Coriell (1961), that performance declined rapidly as cue delay increased. Now, if subjects performing this task apply the cue to the contents of a decaying pre-categorical sensory store, one would expect the declining performance to be caused by a decline in the fidelity with which the cued letter was represented in this store. Reliance on partial visual information should cause item misidentification (which should become more and more frequent as cue delay increases) rather than item mislocation. Instead, wrong responses are usually other items from the display. In other words, what is being progressively lost during the post-display period is information about where items had been located in the display; information about which items had been there is relatively well preserved. In support of this conclusion, Townsend showed in a further experiment that when the post-cue requested only information about item identity, not item location (the cue was a single printed letter and the subject's task was simply to say whether it had been present in the display or not) performance did not decline as cue delay was varied from 0 ms to 450 ms.

These results were confirmed by Mewhort *et al.* (1981), who also showed that misidentification errors (reporting a letter that was nowhere in the display) and mislocation errors (reporting a letter which was in the display but not in the cued position) are differently affected by several variables. Variations in approximation to English affected misidentification errors but not mislocation errors. With a cue delay of 150 ms or more, introducing a backward mask increased mislocation errors but had no effect on misidentification errors. Misidentification error rate was a U-shaped function of retinal locus, whereas mislocation error rate was an M-shaped function of retinal locus.

Mewhort *et al.* therefore argued that the two kinds of errors arise at two different stages. Misidentification errors reflect inadequate visual-feature representation in a pre-categorical sensory-memory stage (which they refer to by the term feature buffer). Mislocation errors reflect inadequate representation of spatial location in a post-categorical stage (a stage they refer to as the character buffer). It was explicitly proposed that the post-display spatial cue is applied to control the transfer from the post-categorical stage to a subsequent output stage, *not* to select items for transfer from pre-categorical sensory memory to the post-categorical stage. Indeed, Mewhort *et al.* proposed that transfer from iconic memory to the post-categorical stage is not selective at all: it is a parallel process (character identification), applied simultaneously to all the items in iconic memory.

Their view is, of course, very much that of Duncan (which has already been discussed) and of others such as van der Heijden (1981): selection operates to transfer material from an initial post-categorical stage to an output stage, *not* to control input to the post-categorical stage. However, there is an important difference between the kind of experiment that Duncan used to support this conclusion and the kind of experiment used by Mewhort *et al.* for the same purpose. One experiment involved selection based on post-categorical information; the other involved selection based on pre-categorical information. This raises the difficult question of how these two types of information can be coordinated.

How is the coordination of episodic-memory and semantic-memory information achieved?

Duncan was concerned with selection based upon the letter–digit distinction. When we speak of items having a post-categorical representation we mean that information about the items that is permanently stored in semantic memory has been accessed – information, for example, about whether an item is a letter or a digit. Thus the mere fact of attaining a pre-categorical representation will mean that information about whether an item is a letter or a digit will be available.

Mewhort *et al.*, on the other hand, studied selection based upon *episodic* information: information about where an item had been located in a particular display. Retrieval from permanent memory ('semantic memory') of everything one knows about an item – that is, categorization of the item – can tell the subject nothing about where the item had been in a multi-item display. If the post-categorical stage contains such information, then representing an item in this stage must involve the coordination of semantic memory and episodic memory to create a representation containing both forms of information. It is by no means clear how this might be accomplished. Consider as an example the pair of words (STEEL mouse):

> Semantic memory tells us that one of (STEEL mouse) is a metal and one is an animal. Episodic memory tells us that what we saw was a word in capitals and one in lower case. What tells us that the metal was printed in capitals and the animal in lower case?

> Presumably, in some way impossible to envisage at present, episodic information becomes temporarily attached to a lexical entry. If this temporary episodic information is subject to rapid decay when it has been obtained hurriedly from a brief display, then perhaps iconic memory consists of the attachment of episodic memory to a lexical entry (Coltheart 1980, p. 67).

The metaphor of attachment to a lexical entry has been unproductive in attempts to explore this question of how the coordination of episodic-memory information and semantic-memory information can occur. More promising, as a way of thinking about this question, are some ideas proposed by Kahneman & Treisman (1983):

> We think of the perceptual system as opening an object file when an object or event is first sensed. The initial entry and the identifying label for the file simply state the location and time. As information about the features of the object is received, it is entered in appropriate slots in the file. Color, size, shape, brightness and direction of movement are specified early, but can be up-dated if and when they change. At some stage, the object may be identified by matching it to specifications in long-term perceptual memory. This allows retrieval and storage in the file of a name or category and of previously learned facts relating to the object, and may also guide the accumulation of further sensory information.

This, it seems to me, is a very promising way of conceiving what I have been referring to as the initial post-categorical stage: the character buffer of Mewhort *et al.* (1981) or the first-level representation of Duncan (1980, 1982). The representation of any item in this form of storage is achieved by creating a temporary file of information about the item. Some of this information concerns sensory attributes of the stimuli as they were presented in a particular display (episodic information) and some is retrieved from the representation of the items in permanent memory (semantic information).

Although Kahneman & Treisman do not discuss the question of a pre-categorical memory stage preceding the file creation stage, such a prior stage must be postulated. The creation of a temporary file must take time, and during this time sensory information must be continuously available. Because we can identify objects even when they are displayed for less than a millisecond, there must be a processing stage which can hold sensory information after the offset of a brief display and can serve as a source of data for the creation of object files. This would be a pre-categorical sensory-memory stage, which one could refer to as iconic memory.

The question of a stage *subsequent* to the object file stage, into which information must be transferred before responses can be made, is less straightforward. However, the data and the arguments of Duncan (1980, 1982) discussed earlier remain apposite, and support the idea that, even after the creation of temporary files representing each one of the items in a multi-item display, if a response is to be made making use of any of the information in one or more of these files, the relevant files must first be selected and transferred to an output stage. This selection procedure is limited in capacity; furthermore,the partial report cue may be thought of as directing this post-categorical selection procedure.

To think of iconic memory experiments as involving the temporary attachment of episodic information to permanent pre-existing lexical entries leaves one unable to say anything about those iconic memory experiments using stimuli that do not possess lexical entries: for example,

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

the experiments of Demkiw & Michaels (1976) and Treisman *et al.* (1975) dealing with iconic memory for moving stimuli. These experiments show that, in multi-item displays in which each item is a meaningless moving stimulus such as a dot of light moving clockwise or anticlockwise, with the subject's task being to report the direction of movement of all stimuli (full report) or a subset of stimuli (partial report), characteristic partial-report superiorities that diminish over time are observed. These cannot be interpreted in terms of the decay of episodic information temporarily attached to lexical entries, because the stimuli have no lexical entries.

The superiority of the Kahneman–Treisman metaphor is marked here. The non-existence of lexical entries is no obstacle to the creation of temporary files representing individual moving items: this non-existence merely entails the absence of lexical information from such files, which can still contain perceptual information such as spatial location and direction of movement. Hence one can think of subjects in the experiments of Demkiw & Michaels and Treisman *et al.* as selecting files in terms of their spatial-location information (because the partial-report cue requested reports concerning items in particular spatial locations), transferring these to the output stage, and making reports based on the information on direction of movement in each transferred file.

Conclusions

As discussed by Gardiner (this symposium), by Coltheart (1983), and by others, there are serious doubts at present concerning whether there exists for auditory material a form of memory that is pre-categorical and purely sensory. The existence of a pre-categorical sensory memory for visual material (that is, of iconic memory), however, is still generally accepted, and has not been challenged here.

The processing of brief visual displays has in the past been explained by theories involving two components, iconic memory (a high-capacity, fast-decaying, pre-categorical visual memory) and some more durable memory mode (low in capacity, post-categorical and with negligible decay).

Certain modifications and extensions of this theory have been advocated here. It is suggested that there exists a third component, a response buffer or output stage, and that items stored in the post-categorical store must first be transferred to the output stage before responses to or decisions about them can be made. This transfer is severely limited in capacity, so that even if many items attain the post-categorical stage, only a few of them can gain access to the output stage.

The post-display cue used in partial report experiments is considered to control the transfer from the post-categorical stage to the output stage. Since subjects can perform the partial-report task when the items to be reported are defined by some pre-categorical (i.e. episodic) feature such as spatial location, colour or brightness, it must be that when items are represented in the post-categorical stage their pre-categorical features remain accessible. As a way of thinking about how this might be accomplished, a metaphor due to Kahneman & Treisman is adopted. The storage of information in the post-categorical stage is viewed as depending upon the creation of temporary item files, each file containing a collection of pieces of pre-categorical and post-categorical information about a particular item that had been in iconic memory. Any such piece of information can be addressed by a partial report cue to cause transfer of the cued item to the output stage for subsequent decisions or for report.

References

Allport, D. A. 1977 On knowing the meanings of words we are unable to report: the effects of visual masking. In *Attention and performance*, vol. 6 (ed. S. Dornic), pp. 505–533: Hillsdale, New Jersey: Erlbaum Press.

Averbach, E. & Coriell, A. S. 1961 Short-term memory in vision. *Bell Syst. tech. J.* **40**, 309–328.

Coltheart, M. 1972 Visual information-processing. In *New horizons in psychology* (ed. P. C. Dodwell), pp. 62–85. Harmondsworth: Penguin Books.

Coltheart, M. 1977 Contemporary models of the cognitive processes. I. Information input and storage. In *The development of cognitive processes* (ed. V. Hamilton & M. D. Vernon), pp. 11–41. London: Academic Press.

Coltheart, M. 1980 The persistences of vision. *Phil. Trans. R. Soc. Lond.* B. **290**, 57–69.

Coltheart, M. 1983 Sensory memories. In *Attention and performance*, vol. 10 (ed. H. Bouma & D. Bouwhuis). London: Lawrence Erlbaum Associates. (In the press.)

Cumming, G. D. & Coltheart, M. 1969 Short-term memory: are item and position information stored independently? *Psychonom Sci.* **15**, 216–217.

Demkiw, P. & Michaels, C. 1976 Motion information in iconic memory. *Acta psychol.* **40**, 257–274.

Dick, A. O. 1974 Iconic memory and its relation to perceptual processing and other memory mechanisms. *Percept. Psychophys.* **16**, 575–596.

Duncan, J. 1980 The locus of interference in the perception of simultaneous stimuli. *Psychol. Rev.* **87**, 272–300.

Duncan, J. 1982 Perceptual selection based on alphanumeric class: evidence from partial reports. Paper presented at Experimental Psychology Society Meeting, Cambridge.

Eriksen, B. A. & Eriksen, C. W. 1974 Effects of noise letters upon the identification of a target letter in a non-search task. *Percept. Psychophys.* **16**, 143–149.

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

Kahneman, D. & Treisman, A. M. 1983 Changing views of attention and automaticity. In *Varieties of attention* (ed. R. Parasuraman, R. Davies & J. Beatty). New York: Academic Press. (In the press.)

Mewhort, D. J. K., Campbell, A. J., Marchetti, F. M. & Campbell, J. I. D. 1981 Identification, localisation, and 'iconic memory': an evaluation of the bar-probe task. *Memory Cogn.* **9**, 50–67.

Sperling, G. 1960 The information available in brief visual presentations. *Psychol. Monogr.* **74**, no. 11.

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

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

Townsend, V. M. 1973 Loss of spatial and identity information following a tachistoscopic exposure. *J. exp. Psychol.* **98**, 113–118.

Treisman, A. M., Russell, R. & Green, J. 1975 Brief visual storage of shape and movement. In *Attention and performance*, vol. 5 (ed. P. M. A. Rabbitt & S. Dornic), pp. 699–721. London: Academic Press.

van der Heijden, A. H. C. 1978 Short-term visual information forgetting. Ph.D. thesis, University of Leiden.

van der Heijden, A. H. C. 1981 *Short-term visual information forgetting*. London: Routledge & Kegan Paul.

Discussion

D. R. J. Laming (*Department of Experimental Psychology, University of Cambridge, U.K.*). In Sperling's experiments the matrix of letters to be reported was preceded and followed by bright pre- and post-exposure fields, equal in brightness to the white of the stimulus. What happens if the matrix of letters is presented in darkness?

What does Professor Coltheart say to this view of iconic memory, an analogy with an electronic filter which presents at output, for subsequent perceptual analysis, a continuous weighted moving average over the preceding input?

D. A. Routh (*Department of Psychology, University of Bristol, U.K.*). About a decade ago there was an emerging controversy about the generality of the interaction between type of cueing (whole or partial report) and length of cue delay. Evidence began to appear that suggested that when whole and partial report trials are randomly mixed within an experiment, as opposed to the blocked arrangement used by Sperling (1960) and others, there is little sign of any convergence of performance for whole and partial report as a function of cue delay (Dick 1971, 1974). Can Professor Coltheart, say, please, what is the present status of this controversy? I have in mind of course the possibility that theorists in this area will have to circumvent yet another instance of a range of context effect (cf. Poulton 1975).

References

Dick, A. O. 1971 On the problem of selection in short-term visual (iconic) memory. *Can. J. Psychol.* **25**, 250–263.
Dick, A. O. 1974 Iconic memory and its relation to perceptual processing and other memory mechanisms. *Percept. Psychophys.* **16**, 575–596.
Poulton, E. C. 1975 Range effects in experiments on people. *Am. J. Psychol.* **88**, 3–32.

D. E. Broadbent, F.R.S. (*Department of Experimental Psychology, University of Oxford, U.K.*). May I put in my usual point, expressing doubt about the possibility that semantic information may be available at a pre-selective stage? Experiments can only show this if they are controlled for differences in the physical features discriminating one category of items from another, and few are. If for instance the task is to detect digits among letters, or vice versa, one must not use curved line digits and straight line letters.

M. Coltheart. Dr Broadbent's point is correct, but fortunately has recently been taken into consideration by Duncan (1982). Duncan compared the detection of digits among letters in two conditions. The first condition used the digits 134568 and the letters CJPQXY; these twelve stimuli were selected so that no visual feature exists that was possessed by all the digits and none of the letters, nor any feature possessed by all the letters and no digit. Duncan's second condition used the digits 235678 and the letters EJINPU; here all the letters have a complete or almost complete vertical contour, whereas no digit has such a feature. Subjects in the first condition, who could not use any physical feature to select digits, yielded degrees of partial report superiority equivalent to those of subjects in the second condition, who could have used physical features for selection.

Dr Routh asks whether, when whole-report and partial-report trials are randomly mixed within an experiment, one can still observe a partial report superiority which diminishes with increasing cue delay. The answer is yes. Among those who have obtained this result are Treisman *et al.* (1975), Merikle (1980) and von Wright (1972).

Concerning Dr Laming's question: it is not true that in Sperling's experiments the letter display was preceded and followed by bright pre- and post-exposure fields. These fields were in fact normally 'dim' or 'dark', according to Sperling (1960), except in one experiment (experiment 5) where it was shown that both full-report and partial-report performances were worse with a bright post-exposure field than with a dark post-exposure field. Averbach & Sperling (1961), however, found that post-exposure luminance affected only partial report, without affecting full report. One would need to resolve this conflict of results before embarking upon consideration of the possible analogy between iconic memory and an electronic filter, although the concept of iconic memory as an integrator of information over time has already been briefly explored by Vanthoor & Eijkman (1973).

References

Averbach, E. & Sperling, G. 1961 Short term storage of information in vision. In *Symposium on Information Theory* (ed. C. Cherry), pp. 196–211. London: Butterworth.
Merikle, P. M. 1980 Selection from visual persistence by perceptual groups and category membership. *J. exp. Psychol.: gen.* **109**, 279–295.
von Wright, J. M. 1972 On the problem of selection in iconic memory. *Scand. J. Psychol.* **13**, 159–171.
Vanthoor, F. L. J. & Eijkman, E. G. J. 1973 Time course of the iconic memory signal. *Acta psychol.* **37**, 79–85.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES