Acta Psychologica 90 (1995) 245-260

acta
psychologica

Does perceptual analysis continue during selection and production of a speeded response?

Jonathan Levy, Harold Pashler *

Department of Psychology 0109, University of California San Diego, La Jolla, CA 92093, USA

Abstract

This study asked whether perceptual analysis of a stimulus can continue while a response to this stimulus is being generated. In Experiment 1, subjects rapidly named a word that was visually degraded with superimposed pixels. Near the time of response, degradation was removed for a short time and then followed by a mask. Subjects then made a second, unspeeded judgment about the identity of the word. Unspeeded judgments were more accurate, showing that the degradation-free stimulus exposure was processed. In Experiment 2, the task was the same, but the degradation was gradually faded out for an individually adjusted duration. Comparison of unspeeded-response accuracy on trials with and without a speeded response showed that stimulus analysis continued at full efficiency during speeded-response generation. The results support conclusions of Rabbitt and Vyas (1981) that perceptual analysis continues during response stages. This form of continuous processing does not necessarily contradict discrete stage models of human information processing, however.

1. Introduction

By its very nature, a speeded choice reaction time task requires a person to accomplish a number of different mental operations. Some of these operations are logically contingent upon others. For example, selecting a response is logically contingent on computing the identity of the stimulus. Beginning with the pioneering work of Donders (1868), it has been suggested that these operations may be temporally discrete, so that a choice task can be decomposed into a succession of stages. If this is the case, then the total reaction time will equal the sum of the durations of each of the component stages. Sternberg (1969) showed that a precise formulation of this kind of stage model makes distinctive predictions for experi-

*Corresponding author. E-mail: hpashler@ucsd.edu.

SSDI 0001-6918(95)00025-9

ments in which factors affecting different processing stages are jointly manipulated (the so-called Additive Factors Method). For the most part, experiments utilizing the Additive Factors method have produced results that are consistent and interpretable in a sensible way (Sanders, 1980).

Nonetheless, many researchers have argued that *continuous flow* models may better describe the actual mode of information transmission in the brain (e.g., Requin et al., 1993). According to these alternative models, the machinery that carries out some particular operation transmits partial information to subsequent, logically contingent stages on a continuous basis, thereby allowing processing to begin at one stage before it has been completed at earlier stages (McClelland, 1979). The disagreement between discrete and continuous views of stages has spawned a substantial empirical literature. Most of this research has focused on the boundary between perceptual analysis and response selection, and has asked whether response selection can begin on the basis of preliminary output from perceptual analysis (see Miller, 1988, for review). A number of experiments by Miller (1982, 1983) suggests that preliminary information does not generally become available (although Miller's data do suggest that when an individual stimulus code - such as color - has been resolved, information may be passed on to response stages before other stimulus codes - such as letter identity - become available).

The studies just described ask whether partial information is transmitted to response-related stages before these latter stages have received the "final" output from the perceptual analysis. The present research addresses a closely related but slightly different issue: whether the perceptual analysis of the stimulus continues at full efficiency even *after* the response selection stage has received all the information that it will use for the generation of an upcoming response. That is, once the "final" decision about the stimulus has been made for the purpose of a given speeded response, does the perceptual system continue to analyze the stimulus along the very same dimension that was relevant in that task? If so, is it able to revise its initial decision without any loss of efficiency due to having come to a preliminary decision? Henceforth, the hypothesis that stimulus analysis continues at full efficiency during the selection and production of a speeded response to a stimulus will be termed *perceptual continuity.*

The "standard" stage diagram - such as the very simple one shown in Fig. 1A - appears to deny perceptual continuity. Indeed, most people who draw such diagrams probably interpret them in that way. Logically speaking, however, perceptual continuity is independent of the question of partial information transmission. For one thing, evidence against partial information transmission (as in Miller's studies) does not necessarily imply that perceptual analysis stops when response selection commences. It is conceivable, for example, that the perceptual machinery might forward partial information about a stimulus, but having arrived at a final decision, analysis of the stimulus might cease (or the system might be "stuck" on whatever decision it came to initially). Conversely, the perceptual machinery might contribute only a single packet of information to the selection of a given response, but nonetheless continue to analyze a stimulus throughout the

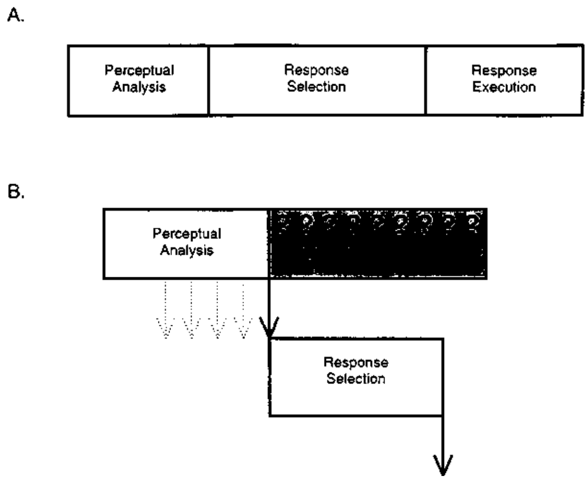

Fig. 1. Panel A shows a simple conventional stage model. Panel B represents the distinction between the two issues discussed in the text: *partial perceptual information* (represented by dotted arrows flowing from perceptual analysis to response selection) and perceptual continuity (represented by ?????)

time this response is being selected and produced. Fig. 1B represents graphically the distinction between the issues of partial information transmission and perceptual continuity. It should also be noted that the validity of the Additive Factor method does not hinge on a denial of perceptual continuity: the assumptions of selective influence and additivity are concerned only with the production of a speeded response, and are indifferent to whether the work of individual stages continues after it no longer matters to the speeded response.

1.1. Evidence regarding perceptual continuity

There are a number of empirical and theoretical considerations that might seem to argue for or against perceptual continuity. When scrutinized, however, none of these turns out to be very compelling. One line of evidence concerns the time course of event-related brain potentials (ERPs) elicited when subjects perform choice-RT tasks. The P300 component of the ERP, which has been linked to perceptual analysis, often peaks during or after response (Duncan-Johnson, 1981). This might seem to indicate that perceptual analysis can continue during response-related stages. The linkage between P300 and percept, however, is not tight enough to justify this argument. P300 latency is increased by factors affecting stimulus analysis but not response selection (McCarthy and Donchin, 1981). This indeed implies that the processes it reflects are dependent upon perceptual analysis. However, it does not demonstrate that the magnitude of the P300 at any given instant measures the "amount" of perceptual analysis going on at that moment. For example, the neural effects generating the P300 might build up even after their trigger was gone, like flooding after a storm. Alternatively, the P300 might reflect processes "downstream" from perception rather than perception itself, such as an affective response triggered by the stimulus analysis. Absent a clear understanding of what instantaneous P300 magnitude means, the fact that it can peak after a response does not provide strong evidence that perceptual analysis of a given stimulus attribute proceeds during generation of a speeded response to that attribute.

Another fact that might seem to argue in favor of perceptual continuity is that perceptual and response-related activity are largely carried out by anatomically distinct brain areas (but see Requin et al., 1993). Why, one might ask, would perceptual areas stop working on their current input when response selection gets underway, if the two types of processes are carried out in separate, "dedicated" modules? This line of argument is not as compelling as it sounds. Even if there were only a one-way flow of information from one set of perceptual machinery to separate response-related machinery, formulating a perceptual decision might still render the perceptual machinery refractory to further input. Indeed, it might be especially refractory to input that would overturn its just-completed decision. (A refusal to revise already-made decisions seems common enough in human behavior viewed at longer time scales.) Indeed, some neural network models of perceptual decision making would seem to make this prediction. For example, in Mozer's (1991) "pull-out network", reaching a decision involves a network settling into a canonical state that would be difficult for subsequent input to overcome.

1.2. Previous research

The most direct approach to the question of perceptual continuity would involve an experiment in which a person makes a speeded response and is then provided a later opportunity to report whatever additional information he or she may have gained from the stimulus. This strategy was first attempted in a very interesting study reported by Rabbitt and Vyas (1981). In one of their experiments, subjects made comparative length judgments of two briefly presented lines which were followed by a mask; the speeded responses were keypresses. Subjects were invited to make corrective responses in the event they realized they made an error. Error corrections did occur, and corrected errors (erroneous responses that were subsequently corrected) were faster than correct responses, which in turn were faster than uncorrected error responses. Rabbitt and Vyas accounted for these results by proposing that perceptual analysis continues during response stages, allowing for the intake of additional perceptual information during response stages. Accordingly, perceptual errors (as opposed to motor errors) could be corrected. On this view, fast incorrect responses benefit from additional perceptual analysis because they occur before the masking pattern; these responses are more likely to be amended. Slow incorrect responses, on the other hand, benefit from less additional perceptual analysis since the masking pattern appears soon after the response; these responses are less likely to be amended.

In order to locate in which processing stage the error occurred, Rabbitt and Vyas (1981) manipulated line length to produce easy and difficult conditions and found that more errors occurred in the difficult condition than the easy one. In the easy condition, they propose that most errors were motor ones. In the difficult condition, they reason that a comparable number of errors were motor ones and deduce that the balance of errors were perceptual ones. These observations are certainly consistent with perceptual continuity, but as Rabbitt and Vyas (1981) acknowledge, they cannot rule out the possibility that corrected errors benefited from additional processing in the response selection or production stages rather than continued perceptual analysis. In addition, the results do not tell us whether continuing perceptual analysis is as efficient as it would be if no early decision were forced, and if the analysis did not overlap with concurrent response selection and generation. Like Rabbitt and Vyas' research, the experiments reported below use a paradigm in which subjects make a speeded response and then make a later, unspeeded judgment on the same stimulus. However, additional features were introduced to help specify the locus of information gain between the speeded and unspeeded judgments, and to assess the efficiency of whatever additional perception was found to be occurring during response-related stages.

2. Experiment 1

This experiment used a naming task. Subjects saw a visually degraded stimulus word and made a speeded response - aiming to produce their response during a particular temporal window *(response window)*. During the later part of the response window interval, the word was briefly undegraded (a *clear flash)* and then masked. Subjects then made a second, unspeeded response. The first response was subject to enough speed pressure to insure that it was based on less than optimal processing time. On many trials, therefore, these responses were incorrect. If perceptual analysis continued during the response stages, additional information should be obtained from the clear flash and result in increased accuracy of the second response.

The predicted improved accuracy of the unspeeded responses would also be consistent with additional processing in response selection and/or response production stage(s). Various control conditions were employed in order to assess these possibilities. Specifically, half of the subjects were not provided with any clear flash. This between-subjects condition will allow us to ascertain whether any improvement evident in the second response accuracy depended on continuing perceptual analysis or on extended processing in response-related stages.

Additionally, all subjects participated in two block types: single- and dual-response. In the latter, subjects made both speeded and unspeeded responses to a stimulus word; in the former, subjects made only an unspeeded response. This within-subjects comparison of unspeeded response accuracy will allow us to assess whether making a speeded response (i.e., decision) affects later processing of the same stimulus. The basic method is shown in Fig. 2.

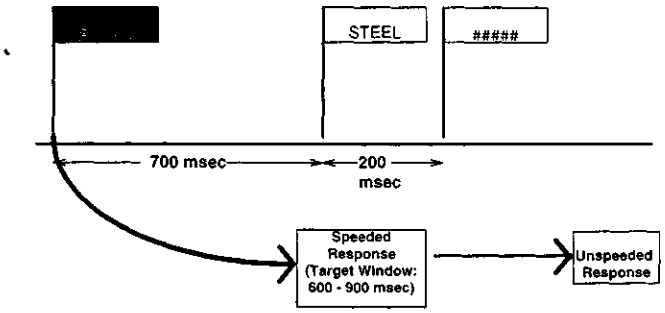

Fig. 2. Paradigm employed in Experiment 1. Subject sees a word degraded by dots (represented by shading in word STEEL). After 700 ms, the word is replaced by a clear flash that lasts for 200 ms, followed by a mask. Subjects make a speeded naming response to the word (aiming for an RT between 600 and 900 ms). Shortly after, they make a "final best guess" about the identity of the word. (In a control condition, there is no speeded response.)

2.1. Method

Subjects

Thirty-two college students participated in one 1-hr session for partial course credit. All subjects reported normal or corrected-to-normal vision.

Design

There were two variables: one within subjects (speeded vs. unspeeded responses) and one between subjects (degraded/clear-flash vs. degraded-only presentation). Half of the subjects participated in the degraded/clear-flash condition and the other half in the degraded-only condition.

A session consisted of ten blocks. For all subjects, there were two types of blocks (single- and dual-response). In the single-response blocks, the subjects made only an unspeeded response; in the dual-response blocks, they made both a speeded and unspeeded response. Each type of block was presented five times, alternating between single-and dual-response blocks. Each block consisted of 30 trials, for a total of 150 trials per block type. Half the subjects for each condition began with the single-response block and the other half with the dual-response block. Stimulus words appeared one time each and were presented in random order.

Apparatus

Stimuli were presented on an NEC Multisync II color monitor, controlled by an IBM PC XT microcomputer (allowing millisecond response timing and synchrony). Subjects made verbal responses into a hand-held microphone connected to a voice key relay (Gerbrandts Model G1341T voice-activated relay), allowing the computer to record the response onset. The verbal responses were also tape recorded.

Stimuli

All stimuli were five-letter nouns with a word frequency rating of ten (according to Kucera and Francis, 1967). The words, composed of blue capital letters, measured approximately 3 cm wide by 1.5 cm high and appeared centrally on a gray background.

In both the degraded/clear-flash and degraded-only presentation conditions, a stimulus word was presented. During the degraded presentation, one thousand white pixels were superimposed in random locations within a central 140 X 40 array measuring 5 cm wide by 4 cm high, thereby covering parts of the word and background. On each trial, the pixel locations were chosen randomly, with replacement. The array, each stimulus word, and each pixel subtended to 4.77° X 3.82°, 2.86° X 1.43°, and 0.096° X 0.034°, respectively, of visual angle from a typical viewing distance of 60 cm.

Procedure

At the beginning of each block, instructions on the screen reminded subjects which block type (single- or dual-response) they were about to begin. For both the degraded/clear-flash and degraded-only presentation conditions, each trial began with the presentation of a central fixation point (a red plus sign) for 500 ms. Next a degraded stimulus word appeared for 700 ms. In the degraded/clear-flash condition, an undegraded (i.e., clear) presentation of the word followed for 200 ms; in the degraded-only condition, there was no clear presentation of the stimulus. In both conditions, a blue masking pattern (# # # # #) immediately followed the word presentation and lasted for 500 ms.

Subjects were instructed that their only task was to name the word. In the single-response block, a single, unspeeded response was made. In the dual-response block, subjects made two responses: a speeded one - with a target response window of 600-900 ms - followed by an unspeeded response. At the end of each trial, subjects were given feedback about the response time on the just-completed speeded response. If the speeded response fell outside the response window, the computer generated a tone and displayed the RT along with a written message indicating whether the response was fast or slow. Subjects made a keypress to initiate the next trial.

Subjects were instructed to try to make all speeded responses within a response window of 600-900 ms, even if this required responding based on incomplete stimulus information in many cases. They were told to make their best guess if they were not sure what the word was by the time they needed to respond. It was emphasized that there was no time pressure for the unspeeded response and that this response could be the same as or different than the speeded response.

The experimenter, who sat in the experimental booth with each subject, recorded the responses with pen and paper at the end of each trial. There was a brief, subject-paced interblock rest period (about 5 s).

All subjects received one practice block of each block type before beginning the session. The words used in practice were not used in the test session.

2.2. *Results and discussion*

Degraded / clear-flash presentation condition

The main results are presented in Table 1. First consider the degraded/clear-flash presentation condition. Here, 77.2% (SD = 7%) of the speeded responses fell within the response window; only these trials were used in the following analysis. The mean accuracy for the speeded and unspeeded responses in the dual-response block was 61.2% (SE = 5%) and 96.7% (SE = 1%), respectively. The difference was significant, $F(1,15) = 67.34$, MSe = 0.01, $p < 0.01$. The mean accuracy for the unspeeded responses, however, did not differ between the single- and dual-response blocks, 97.0% (SE = 0.01) and 96.7% (SE = 0.01), respectively.

The accuracy of unspeeded responses was, not surprisingly, related to the accuracy of the speeded response on the same trial. After a correct speeded response, 99.2% (SE = 1%) of unspeeded responses were accurate, compared to 92.5% (SE = 1%) after an incorrect speeded response. The difference was significant, $F(1,15) = 21.8$, MSe = 0.002, $p < 0.01$. The mean reaction times for correct and incorrect speeded responses were 714 (SE = 4) and 726 (SE = 5) ms, respectively, a significant difference, $F(1,15) = 11.24$, MSe = 102.9, $p < 0.01$.

The result of central interest - improved accuracy of unspeeded responses compared to speeded ones - generally supports the perceptual continuity hypothesis. At least, it implies that some form of additional processing continued to the unspeeded responses. Indeed, even after an incorrect speeded response, a large proportion of the unspeeded responses were correct (92.5%). The fact that unspeeded response accuracy did not differ between the single- and dual-response blocks is consistent with the view that programming the motor response does not adversely affect concurrent perceptual analysis. However, since the data show a ceiling effect, such a claim would be unwarranted.

It is possible that the improved accuracy does reflect continued perceptual analysis, but that the completion of the perceptual analysis (that leads to the speeded response) triggers some degree of refractoriness for additional perceptual analysis. If this were the case, then there would be more time for a second perceptual analysis (following the refractory period) following a fast speeded response than a slow one, given the constant word-mask stimulus onset asynchrony (SOA). Unspeeded response accuracy, therefore, should be higher for trials following a fast speeded response than a slow one. This possibility can be assessed by considering how the accuracy of unspeeded responses relates to the speed of

Table 1
Percentage correct of responses by condition, Experiment 1

Condition	Dual response		Single response
	Speeded	Unspeeded	Unspeeded
Degraded/clear flash	61.2	96.7	97.0
Degraded-only	60.9	75.0	69.9

the speeded responses. Each subject's trials were divided into three bins according to the RT of the speeded response: 600-699, 700-799, and 800-900 ms bins. These contained 45.9%, 37.5%, and 16.7% of the trials, respectively. The mean accuracy for the speeded responses was 64.3% (SE = 5%), 60.6% (SE = 5%), and 54.0% (SE = 5%), respectively, a significant difference, $F(2,30) = 5.40$, MSe = 0.008, $p < 0.01$. Thus, the fastest speeded responses were also more accurate. The mean accuracy for the corresponding unspeeded responses was 97.2% (SE = 1%), 96.7% (SE = 1%), and 96.0% (SE = 1%), respectively. After a correct speeded response, the unspeeded response accuracy was 99.4% (SE = 1%), 99.5% (SE = 1%), and 98.2% (SE = 1%), respectively, and after an incorrect speeded response, it was 93.6% (SE = 2%), 92.5% (SE = 2%), and 91.6% (SE = 3%), respectively. There was no significant effect of bin on these data. Thus, it does not seem likely that perception is suspended for any significant time after the perceptual machinery provides its final input for the speeded response. Interpretation of these results is, however, limited somewhat due to ceiling effects.

Comparison of degraded / clear-flash and degraded-only conditions

Another issue is whether the unspeeded response benefits from specifically perceptual processing, rather than additional processing at some other stage, such as response selection. Between-subjects comparison of the degraded/clear-flash condition just discussed with the control condition - degraded-only - helps in testing this. In the degraded-only condition there was no clear presentation of the stimulus word. Therefore, the only difference between the two presentation conditions is the extent and quality of perceptual input; the time available for response selection processing is the same (i.e., unlimited) in the two conditions. Thus, if extended response selection processing were solely responsible for the improved accuracy of the unspeeded responses, there should be no difference in accuracy of unspeeded responses between conditions. Conversely, to the extent that perceptual continuity is responsible, there should be a big difference.

In the degraded-only presentation condition, the speeded responses of 64.7% (SD = 9%) of the trials fell within the response window; only these trials were used in the analysis. The mean accuracy for the speeded and unspeeded responses in the dual response blocks was 60.9% (SE = 5%) and 75.0% (SE = 4%), respectively, a significant difference, $F(1,15) = 54.07$, MSe = 0.01, $p < 0.01$. The mean accuracy for the unspeeded responses in the single-and dual-response blocks was 69.9% (SE = 3%) and 75.0% (SE = 4%), respectively, a significant difference, $F(1,15) = 12.21$, MSe = 0.002, $p < 0.01$.

After a correct speeded response, 98.7% (SE = 1%) of the unspeeded responses were accurate, compared to 40.3% (SE = 4%) after incorrect speeded responses, a significant difference, $F(1, 15) = 302.3$, MSe = 0.01, $p < 0.01$. The mean reaction times for correct and incorrect responses were 731 (SE = 5) and 746 (SE = 5) ms, respectively, a significant difference, $F(1,15) = 12.59$, MSe = 139.03, $p < 0.01$.

We then compared the degraded/clear-flash and degraded-only conditions. In the dual-response block, the mean accuracy of the speeded responses (61.2% and 60.9%, respectively) showed no significant difference. The mean accuracy of the unspeeded response (96.7% and 75.0%, respectively) did differ significantly, however: $F(1,30) = 30.37$, MSe = 0.01, $p < 0.01$. The difference was even more extreme given an incorrect speeded response: 92.5% in the degraded/clear-flash condition versus 40.3% in the degraded-only condition. This strongly argues that extended perceptual analysis of the clear flash was indeed responsible for the bulk of the improvement.

2.3. Limitations

There are three important limitations to this experiment. First, it is not clear whether continuing perceptual analysis during response stages is as efficient as it would be if there were no speeded decision and concurrent response-related activity. It is possible that it might be substantially less efficient, but we failed to detect this because of the ceiling effect noted above.

The second limitation concerns the nature of the stimulus degradation. The question is whether perceptual analysis of a particular stimulus continues while a response to that stimulus is selected and produced. It was assumed in the preceding discussion that the degraded word and the clear flash constituted a single stimulus. However, it is possible that removing all the degrading pixels at the same time triggers a new round of perceptual analysis that would not otherwise have occurred.

The third limitation concerns the timing of the clear flash. Many speeded responses (those in the 600-699 ms bin) were made before the clear flash. On these trials, then, the perceptual analysis that produced increased accuracy of the unspeeded responses might not have overlapped with the selection and production of the speeded response at all. Furthermore, even when the speeded responses were slow (800-900 ms), there may have been a sizable temporal gap between the completion of the perceptual analysis that fed into the speeded response and the onset of the clear flash.

3. Experiment 2

The second experiment was generally similar to the first experiment, but designed to address the three concerns described above. First, in order to avoid ceiling performance, the stimulus-mask SOA was adjusted for each individual subject in order to obtain 80-85% correct for the unspeeded responses. Second, in order to avoid the problem of the abrupt change introduced by the clear flash, a slightly different degradation technique was employed. All the degrading pixels appeared upon stimulus presentation, but instead of removing them all at once, they were gradually removed, producing the appearance of a fading-in of a single stimulus. For reasons that will emerge below, the experiment also provided some evidence about the efficiency and immediacy of the perceptual processing contributing to the unspeeded response.

In Experiment 1, we compared performance of the degraded/clear-flash and degraded-only conditions in order to ask whether increased accuracy of unspeeded responses reflected specifically continuous perceptual processing. For essentially the same reason, two stimulus-mask SOAs were used within each block in the present experiment.

3.1. Method

The method was the same as Experiment 1 except where noted.

Subjects

Ten new subjects, recruited by poster, participated in four 1-h sessions, one session per day, for pay ($6/session). All reported normal or corrected-to-normal vision.

Design

There were 28 trials per block, resulting in 140 trials of each block type, producing a grand total of 280 trials per session. Half of the subjects began with the single-response block type and the other half with the dual-response block type.

Whereas in Experiment 1, where presentation condition (degraded / clear-flash vs. degraded-only) was a between-subjects variable, in this experiment stimulus presentation was varied on a within-subject basis. There were two stimulus-mask SOA conditions, full and abbreviated, within a block. The abbreviated SOA condition was 200 ms shorter than the full SOA condition. Half of the trials per block were of each type; SOA type varied randomly from trial to trial.

The SOA was adjusted at the completion of every even-numbered block. If the accuracy of the unspeeded responses in the full SOA condition for that block was above 85% correct, then the full SOA was reduced for the next pair of blocks; if the accuracy was below 80% correct, it was increased. SOA remained unchanged for accuracy between 80-85%. All subjects began Session 1 with a full SOA set to 700 ms. Performance on the last block of a session determined the SOA for the first block pair of the following session. Subjects were not informed of the SOA of each trial.

For each session, all subjects were presented with the same group of 280 words in random order. Different words were used in each session.

Stimuli

Stimuli characteristics were the same as in Experiment 1 except the words were of four, five, or six letters; the words measured 2.5, 3.0, or 3.5 cm wide, respectively, by 1.5 cm high.

Pixels were similarly superimposed as in Experiment 1. Immediately upon stimulus presentation, though, 14 randomly selected pixels were removed every refresh cycle (corresponding with the raster scan refresh rate of one cycle per 14 ms). This resulted in a gradual fading-in of the stimulus.

Procedure

Each trial began with the presentation of a central fixation point (a red plus sign) which lasted 500 ms. The stimulus word, with the degrading pixels, appeared next (for either the full or abbreviated SOA), followed by the mask, which lasted 500 ms. After subjects made their verbal response(s) to the stimulus word, they made a keypress which triggered a clear presentation of the stimulus word. They then indicated the accuracy of their verbal response(s) for data collection by making the appropriate keypress(es). (The experimenter was therefore not required to be in the experimental booth to record responses.) In the dual-response block, subjects were next presented RT feedback. After an intertrial interval of a few seconds, the next trial began.

Prior instructions to the subjects stressed the importance of accurate response scoring and informed them that their responses would be verified. They were instructed that if they mistakenly pressed the wrong response key (e.g., pressed "correct" when the response was "incorrect"), they should write down on a piece of paper the word and the appropriate response. Some subjects wrote down no words, while others wrote down only a few *(n* = three to five) words for all four sessions.

At the end of each session, the experimenter reviewed a sample of four to five trials per block of the tape recorded responses to be sure that the subjects were self-scoring with high accuracy. If a discrepancy between responses was detected, the experimenter corrected the response in the data file to match the tape recorded verbal response. For some subjects, no inconsistencies were detected, and for others only a few mistakes (n = three to four) were detected for all four sessions. Thus, subjects were generally very accurate in their self-scoring; although a small proportion of errors were undoubtedly missed, it is very unlikely this could affect the comparisons discussed below.

One practice block of each block type was given before Session 1. In addition, Session 1 itself served as an extended practice session to allow for sufficient SOA adjustment. Data from this session were not used in the analyses.

3.2. Results

The data were collapsed across Sessions 2-4 for each subject. As in Experiment 1, only trials in which the speeded response fell within the response window were used for the analysis. For eight of the 10 subjects overall accuracy was near the target (80-85%) and their SOAs leveled off. Two subjects, however, were below target accuracy over the entire course of the experiment, despite increasing SOAs. We therefore discarded their data before conducting any analyses. Over the course of Sessions 2-4, the average SOA range was 726-811 ms (fastest to slowest SOA). The average difference between fastest and slowest SOA within a subject was 85 ms (SE = 4). The main results are presented in Table 2.

Full presentation condition

On 69.3% (SD = 10%) of the trials, the speeded responses fell within the response window; only these trials were used in the analysis. The mean accuracy

Table 2
Percentage correct of responses by condition, Experiment 2

Condition	Dual response		Single response
	Speeded	Unspeeded	Unspeeded
Full SOA Abbreviated SOA	61.1 57.6	82.3 73.1	79.9 69.8

for the speeded and unspeeded responses in the dual-response block was 61.1% (SE = 3%) and 82.3% (SE = 2%), respectively, a significant difference, $F(1,7) = 23.99$, MSe = 0.023, $p < 0.01$. The mean accuracy for the unspeeded responses in the single- and dual-response blocks was 79.9% (SE = 2%) and 82.3% (SE = 2%), respectively. This difference was not significant.

Abbreviated presentation condition

On 74.8% (SD = 10%) of the trials, the speeded responses fell within the response window; only these trials were used in the analysis. A similar pattern of results were found in the abbreviated presentation condition as in the full presentation condition. The mean accuracy for the speeded and unspeeded responses in the dual-response block was 57.6% (SE = 3%) and 73.1% (SE = 2%), respectively, a significant difference, $F(1,7) = 13.45$, MSe = 0.021, $p < 0.01$. The mean accuracy for the unspeeded responses in the single- and dual-response blocks was 69.8% (SE = 2%) and 73.1% (SE = 2%), respectively. This difference was not significant.

Comparison between full and abbreviated presentation conditions

In the dual-response blocks, accuracy in the full- and abbreviated-presentation conditions differed for both the speeded and unspeeded responses. For the speeded responses, the mean accuracy was 61.7% (SE = 3%) and 57.6% (SE = 3%), respectively, a significant difference, $F(1,7) = 6.0$, MSe = 0.002, $p < 0.04$. For the unspeeded responses, the mean accuracy was 82.3% (SE = 2%) and 73.1% (SE = 2%), respectively, a significant difference, $F(1,7) = 26.16$, MSe = 0.004, $p < 0.001$. The interaction of response and presentation conditions was significant, $F(1,7) = 34.83$, MSe = 0.0002, $p < 0.001$. The mean RT of correct trials was 728 (SE = 6) and 729 (SE = 5) ms, respectively. This difference was not significant.

In the single-response block, the mean accuracy for the unspeeded response in the full- and abbreviated-SOA conditions was 79.9% (SE = 0.02) and 69.8% (SE = 0.02), respectively, a significant difference, $F(1,7) = 29.94$, MSe = 0.004, $p < 0.001$.

3.3. Discussion

The pattern of results in Experiment 2 is similar to those of the previous experiment. Mean accuracy is higher for unspeeded responses than speeded ones for both the full and abbreviated SOA conditions. In the present experiment, comparison of the two SOA conditions allows us to ascertain the locus of this improvement. The mean speeded response RTs did not differ for the full- and abbreviated-SOA conditions, suggesting that subjects performed the task similarly in the two conditions. Speeded response accuracy, however, was slightly higher for the full- than the abbreviated-SOA condition. This implies that the amount of perceptual information used in making the speeded response was slightly less in the abbreviated condition due to the earlier appearance of the mask.

Accuracy was also higher in the full- than in the abbreviated-SOA condition for the unspeeded responses. In addition, the interaction of response and presentation condition was significant: there was substantially more improvement in the full-than the abbreviated-presentation condition. Since both conditions allow equal time for response selection processing, this implies that the improvement in the full SOA condition was largely due to perceptual continuity.

Most critically, there was no sign that unspeeded responses were more accurate in the single-response blocks than in the dual-response blocks - if anything, there was a trend in the opposite direction for both presentation conditions. This argues quite strongly that perceptual analysis was working at essentially full efficiency while the speeded response was being selected and produced. Unlike in Experiment 1, this comparison is not compromised by a ceiling effect.

4. General discussion

In the two experiments reported here, subjects rapidly named visually degraded word. At some point during the stimulus-response interval (or, in some cases, very shortly thereafter), the words were exposed in a less degraded form (abruptly in Experiment 1, gradually in Experiment 2). In both experiments subjects made a second, unspeeded response - their "final best guess" about the identity of the word. By comparing performance in these two conditions (and comparing these against various controls), we attempted to determine whether perceptual analysis continued while the speeded responses were selected and produced. The results of the two experiments support the view that perceptual analysis of a word does indeed continue while a speeded response is being selected and produced. The improved accuracy of unspeeded responses was shown to depend on longer stimulus presentation times in both experiments. Further, the accuracy of unspeeded responses was at least as great in the dual-response blocks as it was when subjects made *only* the unspeeded final judgment (in fact, a trend in the other direction was present). These results imply not only that continued perceptual analysis occurs on the heels of the transmission of a perceptual decision for a speeded response, but also that perceptual analysis continues at or near full efficiency.

4.1. Implications

The most common interpretation of a strict, successive stage model as shown in Fig. 1A is that perceptual analysis terminates once response selection processing

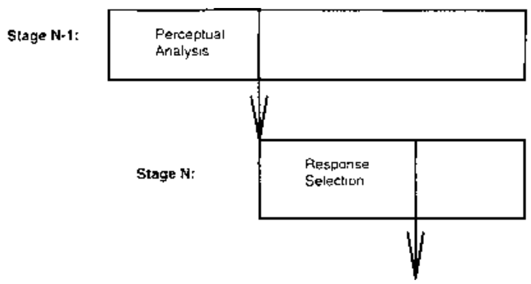

Fig. 3. Discrete stage models (and Additive Factor logic) can be seen as a shorthand for this type of scheme, according to which the initiation of stage *N* (e.g., response selection) is timed to the completion of work on stage *N* - 1 to stage *N*.

commences. The results of this study argue that this is not likely to be correct in general (although, of course, it might be the case with other tasks or stimuli). As noted in the Introduction, however, this conclusion does not challenge the evidence in favor of discrete stage models. Neither, we would argue, does it challenge the validity of the Additive Factor method itself. The fact that the perceptual system continues working on a stimulus does not mean that it necessarily allows information to "trickle in" to the response system before making a final decision. The perceptual decision that triggers a speeded response might depend upon information crossing some distinct threshold, with the response machinery unable to begin work until this threshold is crossed. These data do suggest, though, that diagrams of the discrete stage models like the one shown in Fig. 1A should be interpreted as a shorthand for the more ungainly but apparently more accurate diagram shown in Fig. 3.

How does the hypothesis of perceptual continuity - supported by the present data - fit with findings from other experimental paradigms? It fits well with evidence from dual-task studies suggesting that perceptual analysis of a given stimulus can readily overlap with selection and production of a response to a different stimulus (Pashler, 1989). These studies, too, suggest that perception and response generation can overlap, as one might expect if they are carried out by separate dedicated hardware. It is also interesting to contrast the present results with those situations in which detecting a target of some kind makes it more difficult to detect another target (Duncan, 1980; Shapiro et al., 1994; Pashler, 1994). The two sets of results actually complement each other. The present findings suggest that when perceptual systems transmit information about a given stimulus attribute to response machinery, processing of the object continues along that and other dimensions - in some cases leading to a revision of the decision that was just transmitted. The dual-target effects cited in the preceding sentence suggest, however, that processing of other objects (at least within the same sensory modality) is generally impaired for some time.

Acknowledgements

Support was provided by NIMH (1-R29-MH45584) to H. Pashler. The authors are grateful to Jeff Miller, Allen Osman and Eric Ruthruff for useful discussion.

References

Donders, F.C., 1868. Die Schnelligkeit psychischer Prozesse. Archive der Anatomic und Physiologie, 657-681.

Duncan, J., 1980. The locus of interference in the perception of simultaneous stimuli. Psychological Review 87, 272-300.

Duncan-Johnson, C.C., 1981. P300 latency: A new metric of information processing. Psychophysiology 18, 207-215.

Kucera, H. and W.N. Francis, 1967. Computational analysis of present-day American English. Providence, CT: Brown University Press.

McCarthy, G. and E. Donchin, 1981. A metric for thought: A comparison of P300 latency and reaction time. Science 211, 77-80.

McClelland, J.L., 1979. On the time relations of mental processes: A framework for analyzing processes in cascade. Psychological Review 86, 287-330.

Miller, J.O., 1982. Discrete versus continuous stage models of human information processing: In search of partial output. Journal of Experimental Psychology: Human Perception and Performance 8, 273-296.

Miller, J.O., 1983. Can response preparation begin before stimulus recognition finishes? Journal of Experimental Psychology: Human Perception and Performance 9, 161-182.

Miller, J.O., 1988. Discrete versus continuous stage models of human information processing: Theoretical distinctions and empirical results. Acta Psychologica 67, 1-67.

Pashler, H., 1989. Dissociations and dependencies between speed and accuracy: Evidence for a two-component theory of divided attention in simple tasks. Cognitive Psychology 21, 469-514.

Pashler, H., 1994. Divided attention - Storing and classifying briefly presented objects. Psychonomic Bulletin and Review 1, 115-118.

Rabbitt, P. and S. Vyas, 1981. Processing a display even after you make a response to it. How perceptual errors can be corrected. Quarterly Journal of Experimental Psychology 33A, 223-239.

Requin, J., A. Riehle and J. Seal, 1993. 'Neuronal networks for movement preparation'. In: D.E. Meyer and S. Kornblum (eds.), Attention and performance 14: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience (pp. 745-769). Cambridge, MA: MIT Press.

Sanders, A.F., 1980. 'Stage analysis of reaction processes'. In: G.E. Stelmach and J. Requin (eds.), Tutorials in motor behavior. Amsterdam: North-Holland.

Shapiro, K.L., J.E. Raymond and K.M. Arnell, 1994. Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. Journal of Experimental Psychology: Human Perception and Performance 20, 357-371.

Sternberg, S., 1969. The discovery of processing stages: Extensions of Donders' method. Acta Psychologica 30, 276-315.