

Repetition Blindness has a perceptual locus: Evidence from online processing of targets in RSVP streams

James C. Johnston¹, Larry Hochhaus², and Eric Ruthruff¹

¹NASA - Ames Research Center

²Alcorn State University

The phenomenon of Repetition Blindness (RB) – reduced accuracy in reporting repetitions of briefly displayed items– is often assumed to be a perceptual phenomenon. Several recent studies have found evidence, however, that RB is caused by memory problems. The present experiments measured RB in rapid serial visual presentation (RSVP) streams, with the task altered to minimize memory demands. Experiment 1 freed participants from the need to remember target identities, requiring report of only the total number of targets (one vs. two). Experiment 2 eliminated any biases against repeated targets by segregating repeated and non-repeated targets into separate blocks. Experiments 3 and 4 required immediate online responses to targets as they occurred. All four experiments showed very strong RB. Furthermore, Experiments 3 and 4 provided clear evidence that in RB it is the second of the repeated targets, not the first, that is missed. The present results show that RB occurs “online” as items are initially processed. We argue that RB is indeed a perceptual phenomenon.

In the phenomenon known as “Repetition Blindness” (RB), the accuracy of reporting briefly displayed targets is impaired for repeated targets (Kanwisher, 1987; Marohn & Hochhaus, 1988). Although labeling the phenomenon as a variant of “blindness” implies that it has a perceptual locus, this conclusion was premature. In the traditional paradigm in which RB has been most often investigated, participants view a stream of briefly presented items (Rapid Serial Visual Presentation or RSVP). After the stream of items ends, participants attempt to report the identity of all items. At this point, enough time has elapsed that it would appear to be necessary for participants to retrieve items from memory in order to report them. Hence one must consider the possibility RB is actually a memorial rather than a perceptual phenomenon. This possibility gains additional credibility from the fact that memorial deficiencies in reporting repeated stimuli have been shown to occur even at long exposure durations where perception is

non-problematic (e.g., the “Ranschberg Effect,” Jahnke, 1969, 1972).

Several researchers have, in fact, recently argued for memory retrieval as the processing locus for RB (Armstrong & Mewhort, 1995; Fagot & Pashler, 1995; Whittlesea & Podrouzek, 1995; Whittlesea & Wai, 1997; Whittlesea, Dorken & Podrouzek, 1995; Whittlesea, Podrouzek, Dorken, Williams & Wright, 1995). The basic strategy of these studies has been to hold perceptual encoding factors constant, while manipulating factors that should produce differences in the ease of memory retrieval. Armstrong and Mewhort (1995) and Fagot and Pashler (1995) have shown that the amount of RB found in the delayed-report RSVP paradigm depends on factors that make memory retrieval harder or easier. One particular finding is that RB is strong when participants report all displayed characters (whole report), imposing a severe burden on memory retrieval processes, and weaker when only a subset of items is post-cued to be reported (partial report). There is also evidence that RB in the RSVP paradigm is influenced by response biases against repetitions

operating at the time of memory retrieval and/or overt reporting (Fagot & Pashler, 1995; Whitlesea & Wai, 1997).

These experiments provide strong evidence that in the usual delayed-report RSVP paradigm, RB is modulated by memory retrieval difficulties and by response biases. This evidence does not, however, show that if these influences could be eliminated, then no RB would be found. It remains possible that perceptual processes by themselves are sufficient to produce strong RB.

Other researchers have indeed argued that RB has a perceptual locus (see especially Chun & Cavanaugh, 1997). Based on the finding that RB disappears for items separated by more than 500 ms, Chun (1997) has argued that RB occurs only when perceptual processing of the two stimulus items overlaps in time. Luo and Caramazza (1996) have used a model of detailed parametric data to support a similar conclusion. Both studies concluded that RB is caused by the interaction of perceptual processing of the two targets. Chun and Cavanaugh (1997) found a different way to implicate perceptual processing as a cause of RB. They used apparent motion to control whether two items at a constant distance were or were not perceived as belonging to the same object stream. They found that when apparent motion caused two items to be perceived as states of the same object, RB was much stronger. They argue that the interaction of one clearly perceptual phenomenon, apparent motion, with RB supports the conclusion that RB itself is also a perceptual phenomenon.

This type of evidence is clearly consistent with a perceptual locus for RB. But for each finding of this kind, a determined memory-retrieval theorist could argue that changes in perceptual processing inevitably also affect processing “downstream,” at the time of memory

retrieval, and that it is these downstream memory consequences that modulate RB. Presenting items at longer lags could reduce RB by providing increasingly distinctive associative links to nearby items in the sequence. Putting two identical targets in different object streams could reduce RB by adding distinctive object tags or relative location tags. In general, it is difficult for any stimulus manipulation to be decisive, since that manipulation will inevitably have indirect effects on memory retrieval in addition to direct effects on perception.

We propose a rather different approach to resolving the question of whether RB is a perceptual phenomenon. Our strategy is to refine, not the “front end”, perception-engaging aspects of the RB paradigm, but rather the “back end” memory-retrieval demands. Our goal is to come as near as possible to eliminating any role for offline memory processing. If successful, we should be able to determine whether RB is caused by online perceptual processes that immediately follow target presentation.

Following this strategy, Hochhaus and Johnston (1996) investigated a special “single-frame” paradigm in which participants attempted to identify a single tachistoscopically presented word that either did or did not repeat a preceding supra-threshold prime. Using Signal Detection Theory (Green & Swets, 1966) they found a substantial and statistically significant decrease in sensitivity for repeated targets. This special paradigm is ideal for showing that RB can have a perceptual locus, because there is so little else for the participant to do except encode and immediately report a single stimulus.

The Hochhaus and Johnston (1996) study, however, was only able to create ideal conditions to demonstrate perceptual RB by drastically altering the traditional RB paradigm. A skeptic could still argue that their conclusion has not been shown to generalize beyond the special circumstances of their study. Arguably, making participants responsible only for the contents of a single, difficult-to-perceive item might amount to investigating a different

problem than traditional RB. Note that threshold frame durations were less than half of those used in traditional RSVP RB studies. Perhaps with one very brief frame, the problem with repeated items is perceptual, but with a relatively large number of much longer frames (i.e. in the traditional RSVP paradigm) the problem with repeated items is memorial.

The goal of this paper was to determine if RB is perceptual under the “mainstream” conditions most commonly employed to study RB – RSVP display of many letter identities with frame times usually above 100 ms. In Experiments 1 and 2, participants were required only to report the number of targets found (one or two). This “number-of-targets” task has previously been used by Kanwisher, Kim, and Wickens (1996, Experiment 1). They found that performance was much higher when counting two targets that were two different vowels (e.g., ‘A’ and ‘E’) rather than two occurrences of the same vowel (e.g., ‘A’ and ‘A’). The number-of-targets task has the potential to greatly reduce memory demands because, in principle, non-targets need not be remembered at all. Furthermore only a tally of the number of targets, not their identities, needs to be retained. This procedure also appears to provide little opportunity for contamination of the level of RB by response biases (both repeated and non-repeated target pairs require the same “two” response).

Our description of the potential advantages of the number-of-targets RB paradigm included the words “in principle” because of the implicit assumption that target/nontarget judgments can be made “online”—more or less immediately after display of an item. If the target/non-target classification cannot be made online, then many item identities would still need to be retained until the target/non-target assessment can be made later, offline.

Supposes, for instance, that we declared that, based on sequential position in the alphabet, targets were the odd-numbered letters (e.g. “A”, “C”, “E”, “G”, etc.). If we presented letters at an RSVP display rate of 8 or 10 characters per second, it is wildly implausible that naïve participants could decide online whether “U”, for instance, is or is not a target. This decision would surely have to be made later, offline. In the case of the Kanwisher et al. experiment, participants were precued on each trial with a new pair of vowel targets. It is unclear whether participants actually used the changing precued sets, or simply searched for any vowel. Either way, it is an open question whether or not participants could perform this judgment online. The feasibility of doing so is not as far-fetched as our hypothetical example above, but it cannot be counted on either. With an important theoretical question at stake, further experiments need to be carried out with simpler judgments that are certain to be performed online.

To remedy this problem we used a very simple classification that should be easy to perform online. We used a fixed target set of only two letters: ‘A’ and ‘B’, a set already familiar to participants (the first two letters of the alphabet). Using a fixed set should also ameliorate a problem noted by Fagot and Pashler (1995). They found that when participants had to assimilate a new pair of target letters on every trial, they did much better when the two targets displayed matched the pre-cue sequence (e.g. ‘A’ then ‘B’). This finding raises the possibility that the RB obtained by Kanwisher et al. (1996) resulted not from a generic problem with repeated targets, but rather from this special “congruence bonus.” Using a fixed target set eliminates the pre-cue display that is the source of the problem.

Although we believe that the number-of-targets task combined with a very simple target/non-target judgment reduces memory demands, they are not eliminated entirely, because responses are still delayed until the end of the sequence. Therefore, in Experiments 3

and 4, we attempted to further reduce the role of memory by requiring immediate “online” responses to targets. To preview the results, we found strong RB in all experiments. We believe that this package of experiments represents a considerable advance in efforts to determine whether RB has an “online” perceptual processing locus rather than a memory-retrieval locus. In the General Discussion section we will discuss more generally the merits of distinguishing theoretically between perceptual and memorial processes.

Which of the two repeated targets is missed?

We know that participants in RB experiments often miss one of the repeated targets. But which one do they miss, the first target (T1) or the second target (T2)? According to Kanwisher (1987), RB reflects a failure to individuate the two targets. On this view, it is the second target that is missed (at least, it is not detected as being a new instance of the target type). This prediction is very plausible, but there are other possibilities. For example, it is plausible that display of the repeated target might cause participants to discontinue processing of the first target. According to this account, it might be the first target that is missed.

In a typical RB experiment participants report the identities of T1 and T2. Since T1 and T2 are the same for repeated targets, it is difficult to determine which of the two was missed (see, e.g., Downing & Kanwisher, 1995; Fagot & Pashler, 1995; and Whittlesea & Wai, 1997 for further discussion of this problem). The new paradigm used in experiments 3 and 4, in which participants make speeded responses to each target, provides us with a fresh opportunity to answer this question. In brief, we can tell by the timing of the participant's keypress whether it was likely to have been a response to T1 or T2.

Experiment 1

In Experiment 1 participants judged whether an RSVP letter stream contained one or two targets. To facilitate rapid, online target judgments, the targets on all trials were simply the two letters “A” and “B.” The possible target(s) in any given RSVP stream were therefore ‘A’, ‘B’, ‘AB’, ‘BA’, ‘AA’, and ‘BB.’ Distractor letters were chosen randomly from the remaining 24 letters of the alphabet. Two targets were always separated by one intervening distractor.

Method

Participants: The 12 participants (10 female) were volunteers from among colleges (mean age = 20.3 years). Participants either received course credit or were paid for their participation.

Materials: Testing was carried out in soundproof experimental chambers using PC computers with 21 in. (51 cm) monitors. The stimuli were displayed at a refresh rate of 70 Hz (14.29 ms per refresh cycle). Stimulus characters were white on black in uppercase letters (Genus Font HLV39), displayed one at a time at the center of the CRT screen. At a typical viewing distance of 50 cm, each letter was about 1.45 deg wide and 1.82 deg high.

Procedure: Participants performed a practice block of 16 trials, followed by eight experimental blocks of 34 trials each. Practice trials began at 20 refresh cycles per item (to ensure participants understood the task of reporting the number of distinct instances of A or B targets) and gradually speeding up to only 10 refresh cycles per item.

Each block began with two warm-up trials (not analyzed) followed by 32 trials consisting of four replications of each of eight possible T1/T2 pairs (AX, XA, BX, XB, AA, BB, AB, and BA, where X indicates the absence of a target letter). Filler letters in the RSVP sequences consisted of the 24 non-target letters

with no consecutive repetitions. RSVP sequences varied in length, but on each trial there were always two special item slots where targets could appear, which we will call “slot 1” and “slot 2”. Sequences consisted of 2-5 initial filler items, an item in slot 1, a filler item, an item in slot 2, and then a further 3-6 filler letters. On dual-targets T1 was assigned to slot 1 and T2 was assigned to slot 2. Thus, a sample dual-target sequence might be ‘CJLWAMBOVTN’. On single-target trials, the target letter appeared half the time in slot 1 and half the time in slot 2; a filler letter was assigned to the other slot. Thus the average position of targets on single-target and dual-target trials was equated. Participants pressed one of two keys to indicate whether the RSVP stream contained one or two targets, followed by another keypress to indicate a three-point confidence rating.

Within the 32 test trials of Blocks 2-9, accuracy on the dual-target trials was monitored to determine the presentation rate in the subsequent block. If more than 14 of 16 dual-target trials were correct, the number of screen refresh cycles (14.29 ms) per item was decreased by one. If less than 11 of 16 dual-target trials were correct, the number of refresh cycles per item was increased by one. The mean exposure duration on the final block of the session was 124 ms (range 86 to 172 ms).

In the final data evaluation, both Block 1 (practice) and Block 2 (exposure duration set without feedback from performance) were excluded from analysis. Accuracy was based on the proportion of single-target trials to which the “one” response was made, and the proportion of dual-target trials to which the “two” response was made. The critical comparison for measuring RB was accuracy for repeated vs. non-repeated target pairs within the dual-target condition.

Results

We found a substantial advantage in reporting non-repeated targets (‘AB’ or ‘BA’) compared to repeated targets (‘AA’ or ‘BB’). Mean proportion correct in the dual-target condition was 0.919 for non-repeated targets but only 0.665 for repeated targets, a highly significant effect, $t(11) = 7.44$, $p < 0.001$. Each of the 12 participants had a lower proportion correct for repeated targets ($p < 0.001$ by sign test). Mean proportion correct in the single-target condition was 0.781

Discussion

The data show a very high level of RB-- about four times as high a “miss” rate for repeated target pairs as for non-repeated target pairs-- in a task designed to minimize memory failure and response bias. This outcome extends the evidence for a perceptual locus for RB from the single-frame paradigm of Hochhaus and Johnston (1996) to the more traditional RSVP paradigm. Note that typical exposure durations were well over 100 ms, much longer than the 33-50 ms exposure durations used by Hochhaus and Johnston (1996). Therefore, the present results indicate that very brief exposure durations are not required to produce RB under conditions where memory retrieval problems are minimal.

These results also extend the results of Kanwisher et al. (1996), who used the number-of-targets task, but with a fresh target set (always two vowels) presented before each trial. We discussed earlier the possibility that Kanwisher et al.’s procedure might have made it difficult to perform the judgments online (either because of the workload required to assimilate and use a fresh target set on each trial, or because of the difficulty of making very rapid online vowel/consonant classifications). Furthermore, preview of the target set (e.g., ‘AE’) might have produced a “congruence

bonus” for the corresponding target sequence (e.g., an ‘A’ followed by an ‘E’). In our paradigm, trials do not have a precue, so precue/target congruence problems cannot occur. So the present results show that RB can be obtained in the absence of such congruence effects.

Experiment 2

Upon reflection, it can be seen that the number-of-targets procedure remains open to a different type of bias. Suppose, for whatever reason, participants believed that repeated targets were less likely to occur than non-repeated targets. If so, participants might be biased not to look for repeated stimuli or not to readily accept evidence for repeated targets. In a situation where participants frequently have only partial evidence, such biases could have a large effect.

To deal with this possible bias problem, Experiment 2 again employed the number-of-targets design, but this time repeated targets and non-repeated targets were presented within separate blocks of trials. Participants were told before each block which type of targets were possible for that block and which were not. Note that in the critical block where two targets could only be two repeated targets, a bias against repetitions would make little sense. Thus, it is reasonable to attribute any observed RB effect to the difficulty of detecting repeated targets, rather than a bias against repeated targets.

Method

Except where noted, the method was identical to that of Experiment 1.

Participants: The 16 participants (10 female) were college students (mean age = 21.9 years) tested at the NASA Ames Research Center facility.

Procedure: Participants performed a practice block of eight trials followed by 10 additional blocks of 34 trials each. For one group of eight participants, odd-numbered blocks consisted of an equal number of single-target trials and repeated dual-target trials (AA or BB targets) and even-numbered blocks consisted of an equal number of single-target trials and non-repeated dual-target trials (AB or BA targets). For the other group of eight participants, the conditions of odd-numbered and even-numbered blocks were reversed.

To ensure that average frame duration would be the same for the repeated and non-repeated conditions, frame duration was adjusted only after each pair of blocks (one block with non-repetitions and one block with repetitions). If more than 56 of 64 trials were correct, the number of refresh cycles per frame was decreased by one. If less than 44 of 64 trials were correct, the number of refresh cycles was increased by one. On the final block of the session, the mean exposure duration was 118 ms per item (range: 71 to 157 ms).

Results

We again found a large advantage in reporting pairs of non-repeated targets (‘AB’ or ‘BA’) relative to pairs of repeated targets (‘AA’ or ‘BB’). Mean proportion correct in the dual-target condition was 0.841 for non-repeated pairs, but only 0.578 for repeated pairs, a highly significant difference, $t(15) = 6.71$, $p < 0.001$. Each of the 16 participants showed the same trend ($p < 0.001$ by sign test). Mean proportion correct in the single-target condition was 0.853 (0.844 in blocks with non-repeated targets; : 0.861 in blocks with repeated targets).

Discussion

Experiment 2 again found a strong RB effect that is difficult to attribute to memory retrieval problems. There was no need to store the identity of non-targets, and the only thing

that needed to be retained about targets was a count. These results therefore support the conclusions of Hochhaus and Johnston (1996) and Kanwisher et al. (1996) that RB can occur as an online, perceptual phenomenon.

The present design also further reduces any possible role for response biases in RB. In the number-of-target task, a generic bias against two targets would have affected repeated targets and non-repeated targets equally, without promoting RB. Experiment 2 protected against a more specific bias in favor of non-repeated targets (e.g. a congruence bias in favor of the 'AB' target sequence), or against repeated targets, by blocking the presentation of repeated targets ('AA' and 'BB') and non-repeated targets ('AB' and 'BA'). In blocks where participants know that the only possible dual targets are 'AA' or 'BB', adopting a bias against repetitions would make little sense. (In fact, performance on single targets was almost identical in blocks with repeated targets and in blocks with non-repeated targets, so it appears that differential bias was not an issue.) Because we obtained strong RB with this design, we conclude that bias was not the cause of the RB found in Experiment 1. We believe that the design of Experiment 2 protects against bias rather broadly, including bias in perceptual categorization as well as response bias.

Experiment 3

Experiments 1 and 2 attempted to minimize memory demands by requiring participants to report only the number of target letters in the display. This task clearly makes fewer demands on memory than the traditional whole-report RB task, which requires report of entire sequences of items. However, even the number-of-targets task requires the encoding of some information into memory, and retaining it until the end of the RSVP stream. Furthermore, this

design leaves open the possibility that participants actually memorize as many items as they can, and only compute the count at the end of the trial. Although we consider this possibility remote, further reductions in the role of memory would be welcome.

Experiment 3 used a new paradigm in which participants were instructed to press a response key immediately whenever they detected a target within the RSVP stream. Participants press the same key regardless of which target (A or B) they see. Because participants respond to targets immediately (as opposed to after the entire RSVP stream ends), there is no need to hold any information in memory for more than the time required to initiate a response, and there is no need to retrieve any information from a previous item.

We instructed participants to respond to targets as rapidly as possible. In addition, participants were given a warning message during practice blocks if their response time (RT) to a target exceeded 800 ms. As another measure to discourage reliance on memory, we presented long RSVP streams (usually greater than 20 items) and each stream continued for 1.4 sec following T2. Our procedure was, in fact, quite successful in inducing rapid responding; participants typically responded within about 500 ms of target onsets.

A secondary goal of Experiment 3 was to determine which of the repeated targets is the one being missed, T1 or T2. Traditional RB experiments (and the present Experiments 1 and 2) generally provide information only on how many of the repeated targets were identified. Some experiments have required participants to report entire letter sequences, but even here it is difficult to determine which target was missed (see, e.g., Downing & Kanwisher, 1995; Fagot & Pashler, 1995; and Whittlesea & Wai, 1997). In the present experiment, however, each response comes with a "time-stamp" (i.e., a response time). We can tell from the timing of the response whether it was likely to have been a response to T1 or to T2.

Method

Except where noted, the method was identical to that of Experiment 1.

Participants: The 35 participants (20 female) were college students (mean age = 20.7 years) tested at the NASA Ames Research Center.

Procedure: Participants pressed the ‘J’ key every time they saw a target, regardless of which target was presented (‘A’ or ‘B’). They were instructed to respond immediately to each target, rather than waiting for the RSVP stream to end. Participants were not told how many targets would occur per trial, only that there could be more than one and that the targets could appear very close together in time.

Participants first performed 24 practice trials at exposure durations that began at 20 refresh cycles and declined gradually to 12 cycles. These easy warm-up trials were used to ensure that participants understood the task and knew how to respond appropriately. After this practice sequence, participants then performed 11 additional blocks (consisting of 2 warm-up trials plus 24 test trials). We treated the first three blocks as practice, leaving eight blocks for analysis.

To prevent participants from anticipating targets, the first five items could not contain target slot 1, and for each subsequent item target slot 1 had a 0.2 probability (except that it could occur no later than the 25th item.) The average number of items actually preceding the target 1 slot was 8.87. As in previous experiments, there was always exactly one item between the two target slots

We also increased the number of items following target slot 2. The precise number of items after slot 2 was chosen to ensure that at least 1400 ms elapsed between the onset of slot 2 and the end of the stream (the exact number depended on the exposure

duration). As a consequence, the offset of the RSVP sequence could not be used effectively as a cue that a target had recently occurred, and any response made after that time would have been well over the maximum RT allowed (producing an invalid response).

We again used a staircase to adjust the exposure duration after each block. If performance correctly detected both targets exceeded 0.75 for a block, then the number of refresh cycles per item was decreased by one for the next block. If performance fell below 0.55, then the number of refresh cycles per item was increased by one for the next block. The exposure duration per item on the first experimental block was set at 157 ms. The mean exposure duration was 140 ms on the last block of the experiment (range: 86 to 200 ms).

Results

False Alarms and Misses.

Since this is a detection experiment, only one type of response is made. It is not logically possible to “make the wrong response” per se, but it is still possible to respond at an inappropriate time. Anticipation responses (RT < 100 ms after onset of the first target) occurred on 1% of trials. On single-target trials, participants occasionally made more than one response (6.1 % of trials). When two targets were presented, participants made more than two responses only very rarely (3% of trials). Such low false alarm rates indicate that participants rarely guessed that a target appeared without some supporting perceptual evidence. On single-target trials, participants failed to respond within 1200 ms of target onset on 3.7 % of trials. In the analyses that follow, we count responses on single-target trials as valid if made between 100 ms and 1200 ms after target onset. We count responses on dual-target trials as valid if made between 100 ms after T1 onset and 1200 ms after T2 onset.

Repetition Blindness.

Once again we found a strong RB effect: the probability of correctly detecting both targets was 0.838 for non-repeated pairs and 0.589 for repeated pairs, a highly significant difference, $t(34) = 5.63$, $p < 0.001$. Thirty of the thirty-five participants showed worse performance on repeated targets ($p < 0.001$ by sign test).

Response Time.

In the single-target condition mean RT was 450 ms. In the dual-target condition, mean RT for the first response emitted was 457 ms (447 ms excluding trials where participants responded to only one of the two targets). Thus there was essentially no observed effect of the presence of a second target on RT to the first target. On dual-task trials where at least two responses were emitted, mean RT of the second response was 490 ms (averaged over both repeated and non-repeated target pairs). The 40 ms greater RT to T2 than to T1 is apparently due to a Psychological Refractory Period (PRP) effect, small in size because it is being measured at a stimulus-onset asynchrony (SOA) of about 300 ms (only partial task overlap). Mean second-response RT was similar for repeated (497 ms) and non-repeated targets (482 ms), a difference that did not reach significance, $t(34)=1.74$, $p = 0.092$. Thus, responses to T2 show only a quite small amount of “repetition slowing”.

Discussion

We believe that these results provide the strongest and most direct evidence yet that RB can occur as an online “immediate-response” phenomenon. With a task that required participants to respond immediately to targets as they found them, we observed a very large RB effect—repeating targets doubled the miss rate. These results show that strong RB can be found without the use of an explicit memory-retrieval task. (It

remains true, of course, that in RB the processing of a later T2 is influenced by something left over from processing the preceding T1; so RB still involves some form of implicit memory.)

One clue to the nature of RB is that when participants did detect both targets, responses to a repeated T2 (497 ms) were nearly as fast as responses to a non-repeated T2 (482 ms). This might be viewed as a surprising finding, given the very general tendency for judgments that are more difficult to produce both higher error rates and longer response times. The actual data would clearly be fit rather well by an all-or-none model; either participants failed to detect a repeated target at all, or they detected it in just the same manner as a non-repeated target. In spite of how well this model fits the data, there are several reasons to remain skeptical. Dual-target responses to T2 appear to be subject to a PRP delay, which is often attributed to a bottleneck (cf. Pashler & Johnston, 1989; Welford, 1952). If target repetition alters the rate of processing within a pre-bottleneck stage, then its effect might tend to be reduced by absorption into cognitive slack (see, e.g., McCann & Johnston, 1992; Pashler, 1984; Pashler & Johnston, 1989). Also, there is a possible selection artifact since the analysis excludes trials where no response was made to T2, and the remainder might be a non-random sample. Lastly, we will see that the model does not provide as good a fit to the data from Experiment 4, which showed somewhat more repetition slowing (31 ms). In spite of these cautions, it is still somewhat surprising that the same data set shows such a large effect of repetition on accuracy and yet such a small effect on RT.

Which target is missed, T1 or T2?

It appears that participants in RB experiments often detect only one of two repeated targets, but which one do they miss? In traditional RB experiments there is no reliable way to tell whether T1 or T2 is missed, because

the response to either target would be the same letter identity, emitted at the same time (after the RSVP stream ends). Responses in the current paradigm are emitted on-line and relatively rapidly, however, so we have an opportunity to reliably estimate which target gave rise to each response.

As an overview, our procedure has two phases. In the first phase we determine the RTs of the “extra” responses in the non-repeated dual-target condition that are missing from the repeated-target condition. This is accomplished by simply subtracting the RT histogram for the repeated condition from the RT histogram of the non-repeated condition. In the second phase, we use the resulting difference distribution to estimate how many target misses in the repeated-target condition were T1 misses, and how many were T2 misses.

Figure 1 presents response histograms for the four key conditions of Experiment 3: single-target (T1 only), single-target (T2-only), dual-target repeated, and dual-target non-repeated. In order to facilitate comparisons, the abscissas of the histograms are synchronized to a common starting point (the onset of time-slot 1). The fifth histogram shows the difference between the repeated and non-repeated dual-target histograms.

Single targets always appeared in either time-slot 1 or time-slot 2. Panel A shows the histogram for single targets in slot 1 (i.e. an ‘AX’ or ‘BX’ trial). This histogram has a peak at about 400 ms, and a mean at about 450 ms. The rightmost bar shows the proportion of trials on which the target was missed. Panel B shows the histogram for single-targets in slot 2 (i.e. an ‘XA’ or ‘XB’ trial). The actual RTs to a single-target in the T2 position are no slower or faster, on average, than responses to a single target in the T1 position. But to facilitate comparisons across conditions, in Panel B we bin these RTs relative to the onset of T1.

This amounts to binning the sum of the actual RT to T2 plus the T1/T2 SOA. Panel C shows the histogram for non-repetition dual-target trials (i.e. ‘AB’ or ‘BA’). This is actually a “double-histogram”, binning two responses on each trial; if only one response was made, the other was binned as “missing”. This histogram clearly shows two peaks that line up closely with the peaks for the T1-only (panel A) and T2-only (panel B) conditions. Panel D shows the histogram for dual-target trials with repeated targets (i.e. ‘AA’ or ‘BB’). This histogram also shows two peaks corresponding to responses to T1 alone and T2 alone. Note that the “MISSED” bar is much higher than in the non-repeated condition (panel C), reflecting the RB effect.

The next step is to subtract the repeated histogram from the non-repeated histogram (panel C minus panel D). The results (Panel E) show that the repetition shortfall was concentrated in the bins from about 650 ms to about 1000 ms measured from T1 onset, or about 350 ms to 700 ms relative to T2 onset,

So far we are dealing with empirical facts. How can we best explain them? It is immediately evident: that the missing responses on repeated trials fall in the “second hump” of the bimodal distribution for non-repeated trials (Panel C) and in the range of response times for T2 alone trials (Panel B). The simple hypothesis that explains these facts is that, on repeated-target trials, participants respond to T1 with normal probability and normal RT, but often fail to respond to T2 at all. Hence the RT response that would have been generated to T2 is missing from the distribution. This hypothesis is simple, attractive and provides an extremely close zero-parameter fit to the data. Furthermore, we have not been able to come up with any appealing alternative hypotheses.

Suppose we consider the simplest alternative hypothesis, that on a typical repeated-target trial with only one response, that response was actually made to T2, and the missing response was actually to T1. To explain our results, this hypothesis would require

responses to T2 be extremely rapid, but this could be justified by appealing to priming from T1. There are three reasons why this “missing T1 responses” explanation is implausible: a) Priming would have to be about 300 ms to offset the SOA between T2 and T1, but no repetition priming this large have ever been found with mean unprimed RT’s under 500 ms, b) If the single responses on repetition trials were made in response to T2, their mean RT (relatively to T2 onset) would be about 150 ms and many responses would be under 100 ms—implausibly fast, c) According to this hypothesis, it would be purely coincidence that the single responses made to T2 closely matched the distribution of responses normally made to T1. In contrast, we see no problems with our original hypothesis, described above, that participants tend to miss a repeated T2.

So far, we have resorted only to inspection to confirm that the missing responses to repeated targets (panel E) closely match the distribution of response times to Target 2. In the Appendix, we apply an objective procedure for quantitatively estimating the additional “misses” to T1 and to T2 in the repeated-target condition. The estimate obtained is that repeating targets produced an additional 17 missing responses to T1 (0.5 responses per subject), and 410 missing responses to T2 (11.7 per subject). Given the variability in the error data, the true value for missing T1 responses could easily be zero. As a further quantitative cross-check on our estimates, note that 410 missing T2 responses in a total of 1680 trials is 0.244, accounting very precisely for virtually all of the observed RB (0.589 for repetitions vs. 0.838 for non-repetitions).

In summary, the data from our online task support a surprisingly clear-cut conclusion that RB is almost entirely due to missing responses to repeated second targets.

Experiment 4

Because participants in Experiment 3 responded immediately to targets, the observed RB effect cannot easily be attributed to memory retrieval failure. Instead, the RB effect appears to be due to an online perceptual effect. The purpose of Experiment 4 was to see if this online perceptual effect is due to an inherent difficulty in detecting the repeated item, or to a search bias. We have already discussed the possibility that participants were biased to search for non-repeated target pairs (‘AB’) rather than repeated pairs (‘AA’). To address this issue we used the same technique as in Experiment 2, segregating trials with repeated dual targets and non-repeated dual targets into separate blocks of trials. As in Experiment 3 participants were instructed and trained to respond immediately to targets.

Method

Except where noted, the method was identical to that of Experiment 3.

Participants: The 46 participants (30 female) were college students (mean age = 22.4 years) tested at the NASA Ames Research Center.

Procedure: Participants again pressed the ‘J’ key every time they saw a target, regardless of which target was presented (an ‘A’ or a ‘B’). The primary difference between the present experiment and Experiment 3 is that here we segregated the repeated and non-repeated items into separate blocks of trials (as in experiment 2) Participants completed 12 blocks, alternating back and forth between the two block types. Block-type order was counterbalanced across participants. The first four blocks were considered practice and not analyzed.

Exposure durations were yoked for pairs of blocks with repeated and non-repeated targets. Exposure durations were adjusted after each pair of blocks, aiming for a T2 detection

rate of approximately 65 %. Exposure time per item started the experimental blocks at 157 ms per item, and ended the experiment at a mean of 124 ms per item (range: 71 to 200 ms).

Results

False Alarms and Misses.

Anticipation responses occurred on 2.0 % of trials. When a single target was presented, participants occasionally made more than one response (10.3 % of trials). When two targets were presented, participants rarely made more than two responses (0.8 % of trials). These false alarm rates were modest, considering that each RSVP stream contained a large number of items, each of which provided an opportunity for a false detection. Participants failed to respond to single-targets within 1200 ms on 4.8% of trials (5.0% in blocks where they were mixed with repeated dual targets and 4.5% in blocks where they were mixed with non-repeated dual targets).

Repetition Blindness.

We again observed a strong RB effect: the probability of correctly detecting both targets was 0.681 for repeated pairs and 0.871 for non-repeated pairs, $t(45) = 6.36$, $p < .001$. Thirty-eight of the forty-six participants showed worse performance for repeated items than for non-repeated items ($p < 0.001$ by sign test).

Response Time.

Mean RT to single targets was 439 ms. On dual-target trials RT for the first response was 445 ms (435 ms on trials where two responses were emitted). Mean RT for the second response on dual-target trials was 495 ms. Mean RT to the second target was slightly longer for repeated items (511 ms) than non-

repeated items (480 ms). This 31 ms of “repetition slowing” was significant, $t(48)=5.11$, $p < 0.001$. The interaction across Experiments 3 and 4 of the size of the repetition slowing effect (31 vs 15 ms) was not statistically significant.

Which target is missed, T1 or T2?

Figure 2 shows histograms of responses binned by RT (relative to RT1 onset). The data from Experiment 4 show a pattern very similar to that for Experiment 3. Responses on dual-target trials again showed a bimodal distribution, with the two peaks closely corresponding to the single-trial peaks for targets in slot 1 and slot 2.

As before, panel E is the difference histogram, formed by subtracting the repeated dual-target histogram (panel D) from the non-repeated dual-target histogram (panel C). Panel E again shows that almost all of the missing responses in the repeated condition have RT's similar to T2 RTs. The objective assignment procedure of the Appendix estimates that 31 missing responses were to T1 (0.7 per subject) and 411 missing responses were to T2 (8.9 per subject). The latter estimate corresponds to missing T2 on 0.186 of trials, accounting for virtually all of the observed RB (0.681 repeated targets vs. 0.871 non-repeated targets).

Discussion

The data from Experiment 4 once again showed a strong RB effect. Because participants responded immediately to targets, this RB effect cannot be attributed to memory recall problems. Repeated and non-repeated items appeared in separate blocks of trials, so this RB effect also cannot be attributed to a search bias in favor of non-repeated items. Thus, Experiment 4 supports the conclusion that RB reflects an inherent difficulty in the online detection of repeated targets.

In this experiment, participants responded slightly more slowly when T2 was a repeated item than a non-repeated item.

Although significant, this repetition slowing effect was relatively small (31 ms) in this experiment and even smaller in Experiment 3 (15 ms). These numbers may underestimate the true effect size because of a selection problem-- more repeated-target trials than non-repeated target trials are omitted from this analysis because of missing responses. If real, T2 repetition slowing might provide an interesting clue to the nature of RB. Two alternative possibilities are a) for repeated targets there is refractoriness in the accumulation of stimulus information or in the polling of that information by central processes, and b) information about a repeated T2 accumulates in the same counter as T1, so that a greater increment may be required before detection threshold is reached (cf. the Weber's law model of Hochhaus & Johnston, 1996).

The observed slowing of responses to repeated T2's might provide a means to explore the mechanism or mechanisms underlying RB. If repetition is affecting the very early perceptual encoding operations of T2 processing, then the repetition effect on RT should interact with manipulations of visual quality. Conversely, if repetition is slowing a later operation, such as transfer of detected target to central processes, then the effect of target repetition on RT should be additive with manipulations of visual quality.

General Discussion

Memory demands of RB experiments were reduced using two paradigms. In one paradigm, participants were required only to report at the end of RSVP strings a count of targets, not to their identities. A simple, familiar, and unchanging target set (the letters 'A' and 'B') facilitated counting targets online, as they occurred. Experiment 1, mixing repeated-target trials and non-

repeated-target trials in the same blocks, found very strong RB (0.665 correct reports for repeated targets vs. 0.919 for non-repeated targets). Experiment 2, blocking repeated-target and non-repeated-target trials to reduce any bias against repetitions, also found strong RB (0.578 correct for repeated targets vs. 0.841 correct for non-repeated targets).

In the other paradigm, we reduced memory demands even further by having participants respond online to each target as it occurred. Results showed strong RB in two different experiments. Experiment 3, with repeated and non-repeated targets mixed in the same blocks found strong RB (0.598 correct reports for repeated targets vs. 0.855 correct reports for non-repeated targets). Experiment 4, which segregated repeated and non-repeated dual-target trials into separate blocks, also showed strong RB (0.681 correct reports for repeated targets vs. 0.871 correct reports for non-repeated targets).

Our results significantly extend the generality of previous evidence for strong RB when memory problems are minimized. Hochhaus and Johnston (1996) found strong RB with memory demands at least as small as the current experiments, but their results do not apply straightforwardly to "classic" RB obtained with RSVP stimuli. Kanwisher et al. (1996) previously argued that RB is perceptual based on data from a target-counting task, but it is doubtful whether their task (vowel detection) could be performed online. Our immediate response paradigm confirms directly that our task--searching for the 'AB' target set--can be accomplished online, and that RB can be obtained with immediate overt responses.

Our argument for perceptual RB differs in important ways from previous arguments for perceptual RB. The prevailing strategy has been to show that "front end" manipulations of conditions and circumstances that should influence perceptual processes do in fact modulate RB. For instance, Chun and Cavanaugh (1997) argued for a perceptual locus

for RB because the level of RB is modulated by apparent motion conditions that determine whether two stimuli are seen as a single object. Chun (1997) has argued that the dependence of RB on a relatively short T1-T2 interval also supports a perceptual locus. Similarly, Luo and Caramazza (1996) argued for a specific perceptual cause of RB (refractoriness of type units) using a quantitative model of how RB depends on the overlap in time of activation levels produced by T1 and T2. The problem with such arguments is that almost inevitably any factors that could influence perception would have downstream effects on memory recall as well (cf. the appendix to Fagot and Pashler, 1995). We believe that our strategy—stripping the “back end” off the whole-report RSVP task—avoids this problem, without introducing any equally serious new problems. In any case, the fact that we reach the same conclusion by a different route greatly strengthens the existing case for perceptual RB.

Which Target is Missed?

A bonus of the online response method used in Experiments 3 and 4 is that it provides new and more straightforward evidence about an old RB question: Which target is missed? Because responses were typically made in a relatively narrow time band between 350 and 550 ms, and because multiple targets were presented with an intervening item (resulting in T1/T2 SOAs of about 300 ms) it was possible to assign the great bulk of responses to particular target stimuli (e.g. T1 or T2) with little ambiguity. Analysis of this data shows that almost all of the missing responses to repeated targets compared to non-repeated targets are missing responses to T2. The conclusion that the problem with repeated targets is almost entirely a problem with missing T2, confirms an earlier inference of Park and Kanwisher (1994; cf. also the

discussion in Downing & Kanwisher, 1995). Note also that in the paradigm of Hochhaus and Johnston (1996), only T2 is presented tachistoscopically. Thus in their single-frame paradigm RB is necessarily a problem in perceiving T2. The fact that both paradigms have now been found to implicate a T2 problem is consistent with the conclusion that the two paradigms are studying the same phenomenon, and that RB is generically a problem with the handling of T2.

But is RB “perceptual”?

In this section we tackle head-on a difficult question that we have so far dealt with only implicitly. Does RB deserve to be classified as a “perceptual phenomenon”. Although there is no settled definition of the term “perceptual”, we believe that it is important to sort phenomena into basic categories. Categorization is necessary to determine which clusters of phenomena are good candidates to be explained with common theoretical mechanisms.

The starting point for this discussion is what we know about the nature of RB. The current experiments are very important, because they clearly show that RB is an “online” phenomenon. That conclusion is barely more than a statement of findings from Experiments 3 and 4, where RB was found with immediate speeded responses made shortly after target presentation. We believe it is important that RB can be obtained with typical response times in the range of 350-550 ms, similar to those from other speeded-response tasks with similar modest levels of practice. Thus there is no compelling reason to suspect that anything intervenes between stimulus and response other than the processing stages usually implicated for stimulus classification tasks requiring immediate responses (cf. Sternberg, 1969 and Sanders, 1980).

For our purposes it is useful to decompose the total RT into the stages of stimulus encoding, stimulus classification,

response selection, and response execution. In an experiment where responses are made by pressing a single target key for all targets, we know of no reason why the stages of response selection or response execution should produce more than a trivial error rate; furthermore errors in those stages normally produce commission errors not the omission errors that constitute RB. Hence we conclude that the dramatic increase in miss rates produced by RB must have its genesis in earlier stages. That leaves RB occurring somewhere within the broad stages of stimulus encoding and stimulus classification. Both of these stages correspond well to what is usually meant by “perception”. We do not think it is problematic putting the stimulus classification stage within the “perception” bin. Surely perception includes processes needed to decide whether you have or have not found what you are looking for.

In a nutshell, our argument is that, in a simple immediate-response search experiment, high error rates have nowhere else to come from except “perception”, properly conceived. The detection task is just too simple—and the responses are made too quickly-- for any appreciable proportion of errors to occur anywhere “downstream” of perception. In addition, in an accuracy experiment where errors are induced by reductions in the exposure duration of stimuli (cf. our staircase method), perception should be the natural suspect to be the cause of those errors.

Before provoking unnecessary argument, we should make clear that there is some overlap between the category of “perceptual problems” and the category of “memory problems” conceived most broadly. One common classification of memory problems (e.g. Crowder, 1976) divides them into encoding problems, retention problems, and retrieval problems (where retrieval means simply recall, with

no implication that old experiences can actually be reinstated rather than reconstructed). As far as we can tell, previous researchers who have argued for a memorial cause of RB have clearly been arguing that problems occur during later memory retrieval/recall (which in classical RB experiments occurs after the RSVP string is over, well separated in time from encoding). The RB produced in Experiments 3 and 4 is, on the face of it, incompatible with such a late recall locus¹, because by the time the RSVP string is over, responses have already been made. Our online response procedure also seems to leave little room for retention problems.

On the other hand, we see no reason to oppose including the RB phenomenon studied here in the broad category of encoding problems. There is clearly a close relationship between a) the transfer of information from peripheral perceptual analyzers to central processes capable of initiating responses, and b) the storage of information into memory buffers (cf. Jolicoeur, 1999). Whether these ultimately turn out to be one and the same or only closely intertwined cannot yet be determined. Perceptual processing would appear to be an ingredient in almost all memory encoding, and transfers among memory buffers are commonplace in perceptual processing (especially with learned target categories) .

But there is also a clearly defined part of the memory domain that is not so strongly interwoven with perception, namely retrieval of information after it has ceased to be a part of the immediate present or primary memory. We argue that the present results show that strong RB can be obtained under conditions where it is clearly not such a later, post-encoding, memory-retrieval problem.

Previous evidence that RB is memorial.

In this section we discuss how the present results can be reconciled with previous studies arguing that RB is caused by memory-retrieval failure (Armstrong & Mewhort, 1995;

Fagot & Pashler, 1995; Whittlesea & Wai, 1997; Whittlesea et al, 1995). Most importantly, we are not claiming that all reported RB must have online/perceptual causes. We believe that it has been well demonstrated that various difficulties in memory retrieval can play a substantial role in RB.

Whittlesea and Wai (1997), for instance, have made a convincing case that whole report of sentences—the paradigm in which RB was first documented by Kanwisher (1987)—makes use of a reconstructive process (see Bartlett, 1932) that includes inferences from the syntactic and semantic context. Whittlesea and Wai (1997) find that in this reconstructive process, repetitions are at a disadvantage and are less likely to be retrieved and/or reported. In making a whole report, participants need to keep track of which items in memory have already been reported; repetitions present a special problem because it is difficult to distinguish a second token of the same type from an (undesired) resampling of a single token (cf. Fagot & Pashler, 1995). Using a partial report paradigm, in which a report cue is presented after the RSVP stream is over (e.g. Fagot & Pashler, 1995), does not eliminate all memory problems. Participants still must encode the entire string into memory, and engage in memory retrieval operations after the RSVP stream is over.

In summary, many RB studies have been conducted using paradigms that make important demands on memory, and there is evidence that the RB found in those paradigms is at least modulated by memorial factors.

Evidence that memory manipulations influence the level of RB poses no challenge to our position. We believe that in RB paradigms where the level of accuracy is determined by both perceptual and memorial difficulties, RB will typically have both

perceptual and memorial causes. There remains, however, one puzzle we would like to address. Several papers have made the stronger argument that when memorial problems are eliminated, RB actually disappears (e.g. Experiments 2 and 3 of Armstrong & Mewhort, 1995; Experiments 2, 3 and 4 of Fagot & Pashler, 1995). Since we claim that perceptual processes by themselves typically produce RB, and since perceptual encoding was still involved in these paradigms, should not we expect RB to still have occurred?

In reply, we first note that we are not claiming that perceptual processes inevitably cause RB, only that under appropriate circumstances they can cause RB. While this generic response should make it clear that failures to find RB do not actually contradict our conclusion, one might still argue that they limit the generality of our conclusions. Could it be that, in spite of the present empirical results, perceptual causes do not in general tend to produce RB?

It is useful to look at the specifics of the experiments purporting to show that RB disappears when memorial problems are eliminated. The experiments of Armstrong and Mewhort (1995) failed to find RB using a cueing procedure that required reporting the letter that had followed the cued letter in the RSVP stream. This procedure is most unusual, placing a burden on participants of not just identifying letters but also encoding their sequence. Not surprisingly, performance was very poor (in Experiment 2, accuracy was 0.26 for repetitions vs. 0.25 for non-repetitions). The low level of performance on non-repeated letters provides a poor baseline from which to detect further decrements due to RB. Furthermore, this poor performance raises the question of how well participants could make use of the cue. If they typically did not use the cue—in which case the task amounted mostly to choosing at random a letter they had seen to report—then a subtle artifact is present. For repetitions, there is clearly some probability of picking T1 when

responding to a T2 query, which would count as correct; for non-repetitions, reports of T1 to a T2 query would not count as correct. This “reports of T1 count as reports of T2” artifact would distort the data in favor of repetitions, masking the true level of RB. The authors make some indirect arguments against guessing models, but it is implausible that guessing did not occur, and there is simply no way to avoid counting reports of T1 as correct on repetition trials.

Fagot and Pashler (1995) reported four different experiments showing an absence of RB. This is impressive, but all four shared an unusual aspect—there was some non-standard perceptual difference between T1 and T2. In Experiments 2 and 3, list items were strung out spatially so that T1 and T2 always occurred in different locations. In Experiment 4, the critical data showing an absence of RB come from trials where T2 was colored red and T1 was white (and subjects knew red was a likely retrieval cue). In Experiment 5, some items were visual and some were auditory and repetitions always occurred across modalities.

When T1 and T2 can be distinguished by other prominent ancillary perceptual properties, it is questionable whether any perceptual RB should be expected. Distinguishing properties like location, color or modality should increase the ease with which subjects can collect evidence in separate counters (Hochhaus & Johnston, 1996) or achieve token individuation² (Kanwisher, 1987). This point is directly supported by the finding of Chun (1997) that no RB was found when the two targets were colored red and green in a stream of black items. In summary we would argue that the Fagot and Pashler (1995) experiments do suggest boundaries for the conditions promoting RB, but these boundaries are congenial to existing perceptual RB theories (see next section).

Theories of Repetition Blindness.

The present data are consistent with any theory that attributes RB to online perceptual operations. The data are consistent, therefore, with token individuation theory of RB (Kanwisher, 1987; Chun, 1997). This theory distinguishes between object recognition (making contact with the long-term representation of a visual type) and object individuation (forming a token that localizes that object in time and space). Applied to RB, token individuation theory hypothesizes that in many instances the repeated target in an RSVP stream is recognized, but not individuated. As a result, the participant is sometimes aware of and reports only one instance (token) of the repeated target.

The data are also consistent with a similar theory proposed by Hochhaus and Johnston (1996) to explain the RB effects found in their “single-frame” paradigm. Their participants saw a pre-cue word followed by a tachistoscopic presentation of the target word. Identification of the target was markedly worse when the target matched the pre-cue than when it did not. Hochhaus and Johnston used a generalization of the classical Weber-Fechner law to explain their results. First, assume that there are logogen units, each with a certain level of activation, corresponding to each possible candidate target word. The participants’ task is to determine which logogen unit received the greatest increase in activation following the tachistoscopic display. This judgment becomes very difficult when the target word matches the precue, because the corresponding logogen was already highly activated by the precue. Even if a new target were able to raise the logogen activation level by the usual increment, this would constitute a much smaller Weber ratio of increment to base activation than would occur without the precue. Note this theory is not necessarily incompatible with Kanwisher’s (1987) token individuation theory. The smaller Weber ratio could be the reason why the perceptual system does not create a new token

for the repeated object – the same change detection threshold that is normally passed for non-repetitions would be much less likely to be passed for repetitions.

Concluding Remarks

The present experiments show that RB occurs even when the role of memory demands and response biases are minimized. RB in RSVP streams can be measured a) in a search paradigm in which participants only need to count the number of targets, and b) in an immediate speeded-response paradigm in which participants respond to each target as they detect it. These data provide strong support for the hypothesis that RB has an online perceptual cause rather than a cause at the later stage of memory retrieval. Data from the immediate response paradigm also reveal that in RB it is the response to the second target that is missing. Further work is needed to more precisely determine the locus of the perceptual problem underlying RB.

References

- Armstrong, I. T., & Mewhort, D. J. K. (1995). Repetition deficit in rapid-serial-visual-presentation displays: Encoding failure or retrieval failure? *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 1044-1052.
- Bartlett, F. C. (1992). *Remembering: A study in experimental psychology*. Cambridge, England: Cambridge University Press.
- Chun, M.M. (1997) Types and tokens in visual processing: a double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, **23**, 738-755.
- Chun, M.M. & Cavanaugh, P. (1997). Seeing two as one—linking apparent motion and repetition blindness. *Psychological Science*, **8**, 74-78.
- Crowder, R.G. (1976). *Principles of learning and memory*. Hillsdale, NJ: Erlbaum.
- Downing, P., & Kanwisher, N. (1995) Types and tokens unscathed: A reply to Whittlesea, Dorken, and Podrouzek (1995) and Whittlesea and Podrouzek (1995). *Journal of Experimental Psychology: Learning, Memory and Cognition* **21**, 1698-1792.
- Fagot, C., & Pashler, H. (1995). Repetition blindness: Perception or memory failure? *Journal of Experimental Psychology: Human Perception and Performance*, **21**, 275-292.
- Green, D.M., & Swets, J.A. (1966). *Signal Detection Theory and psychophysics*. New York: Wiley.
- Hochhaus, L., & Johnston, J. C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance*, **22**, 355-366.
- Jahnke, J. C. (1969). The Ranschburg effect. *Psychological Review*, **76**, 592-605.
- Jahnke, J. C. (1972). The effects of intraserial and interserial repetition on recall. *Journal of Verbal Learning and Verbal Behavior*, **11**, 706-716.
- Jolicoeur, P. (1999). Dual-task interference and visual encoding. *Journal of Experimental Psychology: Human Perception and Performance*, **25**, 596-616.
- Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, **27**, 117-143.
- Kanwisher, N. G., Kim, J. W., & Wickens, T. D. (1996). Signal detection analysis of repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, **22**, 1249-1260.
- Luo, C.R. & Caramazza, A. (1996). Temporal and spatial repetition blindness: Effects of presentation mode and repetition lag on the perception of repeated items. *Journal of*

- Experimental Psychology: Human Perception and Performance*, **22**, 95-113.
- Marohn, K. M., & Hochhaus, L. (1988). Semantic priming increases and repetition priming decreases apparent stimulus duration. *Journal of General Psychology*, **115**, 51-61.
- McCann, R.S., & Johnston, J. C. (1992). Locus of the single-channel bottleneck in dual-task interference, *Journal of Experimental Psychology: Human Perception and Performance*, **18**, 471-484.
- Park, J., & Kanwisher, N. (1994). Repetition blindness: Effects of timing, load, and stimulus configuration. *Journal of Experimental Psychology: Human Perception and Performance*, **29**, 508-519.
- Pashler, H. (1984). Processing stages in overlapping tasks: Evidence for a central bottleneck. *Journal of Experimental Psychology: Human Perception and Performance*, **10**, 358-377.
- Pashler, H., & Johnston, J. C. (1989). Chronometric evidence for central postponement in temporally overlapping tasks. *Quarterly Journal of Experimental Psychology*, **41A**, 19-45.
- Sanders, A.F. (1980). Stage analysis of reaction processes. In G.E. Stelmach and J.Requin (eds.) *Tutorials in Motor Behavior*. Amsterdam: North-Holland, pp-331-354.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. In W.G. Koster (ed.), *Attention and Performance II*. Amsterdam: North-Holland, pp. 276-315.
- Welford, A.F. (1952). The "psychological refractory period" and the timing of high-speed performance—A review and a theory. *British Journal of Psychology*, **43**, 2-19.
- Whittlesea, B. W. A., Dorken, M. D., & Podrouzek, K. W. (1995). Repeated events in rapid lists. Part 1: Encoding and representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1670-1688.
- Whittlesea, B. W. A., & Podrouzek, K. W. (1995). Repeated events in rapid lists. Part 2: Remembering repetition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1689-1697.
- Whittlesea, B. W. A., Podrouzek, M. D., Dorken, M. D., Williams, L. D., & Wright, R. L. (1995). Types and tokens: Unscathed or unscathable? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **21**, 1703-1706.
- Whittlesea, B. W. A., & Wai, K. H. (1997). Reverse "repetition blindness" and release from "repetition blindness:" Constructive variations on the "repetition blindness" effect. *Psychological Research*, **60**, 173-182.

Author Note

We thank Beth Siegel and Julie Creagh for their assistance, and we thank the NASA Ames Research Center for their support.

Correspondence concerning this article should be addressed to James C. Johnston, Mailstop 262-4, NASA Ames Research Center, Moffett Field, CA 94035.

Appendix

Estimating the effect of repeated targets on detection of T1 and T2

In both Experiment 3 and Experiment 4 we derive distributions of “missing responses” on repetition trials (Figure 1E and 2E). In each case inspection suggests that these distributions follow closely the distribution for T1 presented alone, and also the distribution of the rightmost hump of the distribution for non-repeated two-target trials. In the text we use these observations to support the hypothesis that repetition trials suffer mainly from misses of T2.

In this appendix we address the question of whether it is possible to establish an objective criterion that can be used to classify the missing responses on repetition trials as T1 misses or T2 misses. One simple idea would be to find the median responses times for the T1 alone and T2 alone distributions and use the mid-point of that range as the dividing line between T1 and T2 misses. (In fact, if this simple “split-the-difference” heuristic is applied to the data, the numbers are very little changed from those we are about to obtain).

This “split-the-difference” procedure will produce classification errors of two kinds. Some responses that actually belong in the “slow” tail of T1 responses will be incorrectly assigned to T2, and some responses actually from the “fast” tail of T2 will be incorrectly assigned to T1. Because RT distributions are almost inevitably positively skewed, it is likely that more of the first type of errors will occur than the second. To minimize classification errors, we decided to set the criterion at the crossover point along the RT axis in the frequency of T1 alone and T2 alone responses. Figure 1E shows that the response-time bin for 525-575 contains substantially more responses in the T1 alone bin than the T2 alone bin (where responses would have been very fast, about 300 ms after T2 onset). One can also see that the next response-time bin for 575-625 ms contains substantially more responses for T2 alone than for T1 alone. Hence the cross-over point for the relative likelihood that a response came from the T1 or T2 distributions is near the 575 ms value that divides these bins. By setting the criterion at 575 ms, we can closely approximate the optimal criterion for separating responses due to T1 and T2. Using this criterion, 17 missing responses (0.5 per participant) are classified as responses to T1, and 410 (11.7 per participant) are classified as responses to T2.

For Experiment 4 the same procedure was used to set the criterion dividing the T1/T2 classification. The data were so similar that the criterion chosen had the same value as for Experiment 3--575 ms. Using this criterion, 31 missing responses (0.5 per participant) are classified as responses to T1, and 411 (8.9 per participant) are classified as responses to T2.

It is possible to argue that the criterion used should be adjusted either higher or lower. The argument for raising the criterion is that the appropriate baseline should not be T2 alone but rather T2 dual-target, for which the times are slightly longer. The maximum adjustment called for would be no more than a few tens of ms (at most one bin). The argument for lowering the criterion is subtle. If one accepts the overall conclusion that there are many more misses of T2 than T1, then we have a classic unequal base-rates problem. A trial right at our criterion would actually be much more likely to be a T2 error. Hence a lower criterion would be needed to offset the base-rate difference. (Note that because of this consideration, even the tiny estimate of T1 misses may be an over-estimate).

Fortunately, nothing of theoretical importance hinges on the exact numbers of the estimates. Looking at the distributions, it is clear that nothing would change the conclusion--that almost all misses on repeated-target trials are T2 misses--except a very large increase in the

criterion, for which there is no motivation. So we believe that any reasonable classification method would reach the same conclusion.

Table 1

Proportion correct responses for non-repeated and repeated dual-target conditions and the size of Repetition Blindness effects in Experiments 1-4

	Design	Non-repeated Dual Targets	Repeated Dual Targets	Repetition Blindness Effect Size
Exp 1	Mixed	.919	.665	.254
Exp 2	Blocked	.841	.578	.263
Exp 3	Mixed	.838	.589	.249
Exp 4	Blocked	.871	.681	.190

Figure Captions

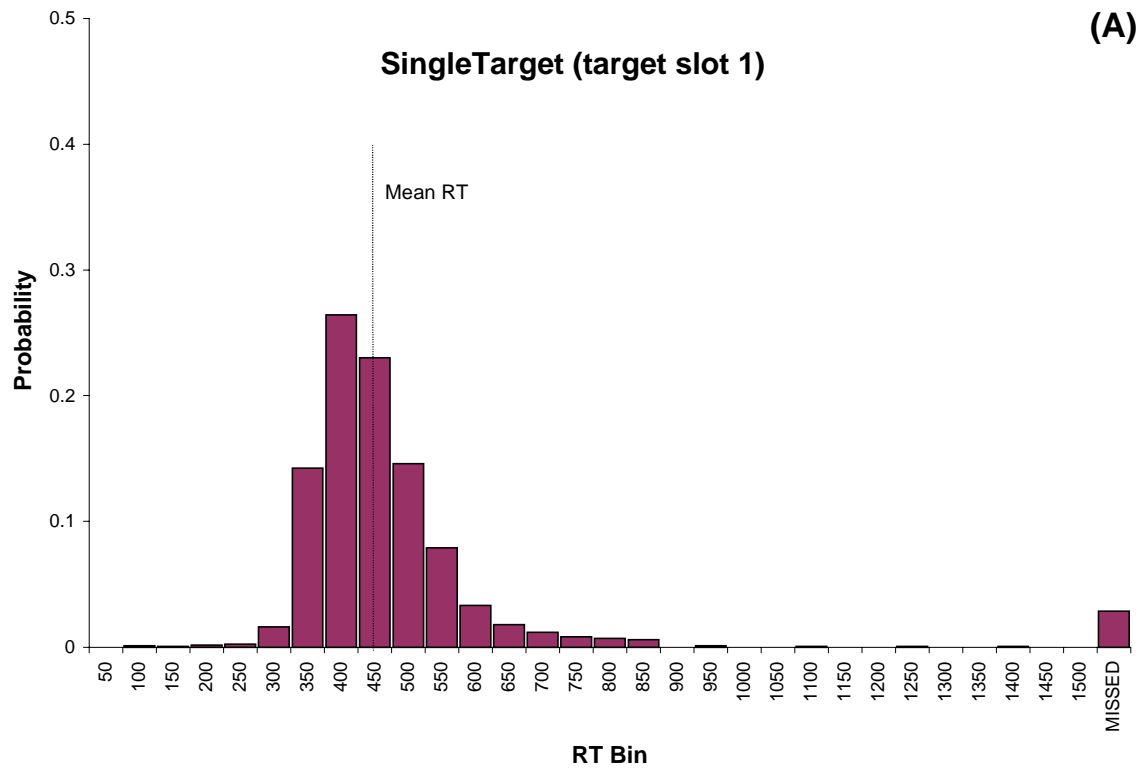
Figure 1. Histograms for the data from Experiment 3. The abscissa represents response-time (RT) bins 50 ms wide, measured from the onset of target slot 1. . The rightmost bar shows the proportion of trials where the participant missed the target (or the RT was greater than 1525 ms). Panel A shows the histogram for single-target trials where only the target slot 1 contained a target. Panel B shows the histogram for single-target trials where only target slot 2 contained a target. Panel C shows the “double” histogram (two responses binned per trial) for dual-target trials without target repetitions. Panel D shows the “double” histogram (two responses binned per trial) for dual-target trials with repeated targets. Panel E shows the histogram resulting from the subtraction of panel D values from panel C values, estimating the RT distribution of the missing responses to repeated targets.


Figure 2. Histograms for the data from Experiment 4. The abscissa represents response-time (RT) bins 50 ms wide, measured from the onset of target slot 1. . The rightmost bar shows the proportion of trials where the participant missed the target (or the RT was greater than 1525 ms). Panel A shows the histogram for single-target trials where only the target slot 1 contained a target. Panel B shows the histogram for single-target trials where only target slot 2 contained a target. Panel C shows the “double” histogram (two responses binned per trial) for dual-target trials without target repetitions. Panel D shows the “double” histogram (two responses binned per trial) for dual-target trials with repeated targets. Panel E shows the histogram resulting from the subtraction of panel D values from panel C values, estimating the RT distribution of the missing responses to repeated targets.


FOOTNOTES


¹ Of course even immediate-response tasks involve passing information from one processing stage to the next, and this may require very brief periods of buffering information, either as output buffering of any stage N or input buffering of stage N+1. We have no problem classifying a phenomenon as perceptual that involves problems during storage into or retrieval out of very short-term perceptual buffers.

² A reviewer pointed out that Kanwisher's original theory of token individuation did not allow for multiple tokens to coexist simultaneously at all. But here we are interested here in the larger set of possible theories that attribute to difficulties with token individuation. It seems clearly possible that a token individuation problem with otherwise identical token stimuli could readily be ameliorated when differential stimulus properties are introduced. This would make it possible to gracefully handle several instances of the same token word, each joined in a bundle with different distinguishers such as color (cf. Chun, 1997).

Figure 1







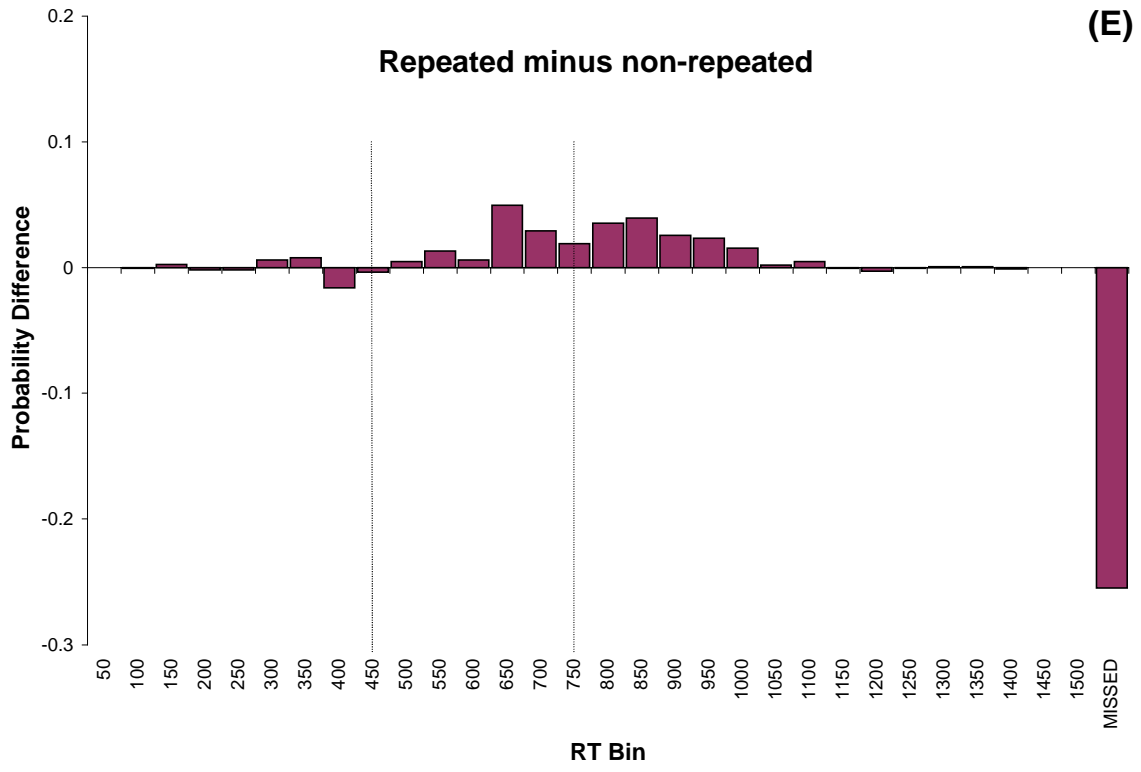


Figure 2