



Repetition blindness for words and pictures: A failure to form stable type representations?

Irina M. Harris¹ · William G. Hayward² · Manuel S. Seet³ · Sally Andrews¹

Accepted: 22 January 2021 / Published online: 5 March 2021
© The Psychonomic Society, Inc. 2021

Abstract

Repetition blindness (RB) is the failure to detect and report a repeated item during rapid serial visual presentation (RSVP). The RB literature reveals consistent and robust RB for word stimuli, but somewhat variable RB effects for pictorial stimuli. We directly compared RB for object pictures and their word labels, using exactly the same procedure in the same participants. Experiment 1 used a large pool of stimuli that only occurred once during the experiment and found significant RB for words, but significant repetition facilitation for pictures. These differential repetition effects were replicated when the task required participants to only report the last item of the stream. Experiment 2 used a small pool of stimuli presented several times throughout the experiment. Significant RB was found for both words and pictures, although it was more pronounced for words. These findings present a challenge to the token individuation hypothesis (Kanwisher, *Cognition*, 27, 117–143, 1987) and suggest that RB is more likely to be due to a difficulty in establishing a robust type representation. We propose that an experimental context that contains high levels of overlap in visual features (e.g., letters in the case of words, visual fragments in the case of repeatedly presented pictures) may prevent the formation of distinct object-level episodic representations, resulting in RB.

Keywords Repetition blindness · Types · Object tokens · Stimulus format

Rapid serial visual presentation (RSVP) of information, in which visual stimuli are presented at rates of 8–12 items/s in the same spatial location, can induce surprising failures of attentional processing and conscious awareness. One such failure is repetition blindness (RB), which occurs when an item is repeated after a short temporal lag (within 500 ms or so). Under these conditions, observers are frequently unaware of the occurrence of the repetition and are much less likely to report a repeated stimulus than a stimulus with a different identity. The original explanation of RB was that, while information about the *type* of the stimulus is registered at some level—accounting for the fact that performance is sensitive to the identity of the items—a second episodic representation, or *object token*, of the repeated item is not established under RSVP conditions because of limitations in the capacity to bind

two tokens to the same type within a short amount of time (Kanwisher, 1987). This *token-individuation* explanation has remained the dominant view, although others have suggested alternative explanations, including response and memory reconstruction biases (Fagot & Pashler, 1995; Whittlesea et al., 1995; Whittlesea & Masson, 2005; Whittlesea & Podrouzek, 1995) and competition for awareness (Morris et al., 2009).

All the above explanations assume that the type representation of the repeated stimulus is accessed, and the difficulty lies in subsequent processes required for this representation to become consciously available. A different account is offered by the *type-refractoriness account* (Luo & Caramazza, 1995, 1996), which assumes that RB arises because the “type node” for a particular item has a refractory period after activation during which it cannot be activated again (Luo & Caramazza, 1995). This is an idea that Kanwisher (1987), who first reported the RB phenomenon, had dismissed on the basis of a finding that RB only occurred when both of the repeated items (usually denoted by C1 and C2, for Critical Items 1 and 2) had to be reported. If observers only had to report the last item in an RSVP list, accuracy was higher when this was a repetition of an earlier item in the list (Kanwisher, 1987, Experiment 3). The fact that priming is seen under these circumstances, instead of RB, would be

✉ Irina M. Harris
irina.harris@sydney.edu.au

¹ School of Psychology, Brennan-MacCallum Building A18, University of Sydney, Sydney, NSW 2006, Australia

² University of Hong Kong, Hong Kong, Hong Kong

³ National University of Singapore, Singapore, Singapore

inconsistent with the idea that the critical type is in a refractory state when C2 occurs. However, Kanwisher's original grounds for rejecting the type refractoriness hypothesis are challenged by a number of failures to replicate the finding that an unattended C1 yields priming rather than RB. These experiments have shown that RB can be obtained even when C1 is not reported (Burt & Jolley, 2017; Kanwisher & Potter, 1990; Leggett et al., 2019; Luo & Caramazza, 1995), leaving open the possibility that RB can be caused by difficulties in establishing a robust type representation. This was examined in the current study.

An additional reason to question the idea that RB is due to token individuation of otherwise robustly activated types is provided by evidence of variations in the strength and extent of RB obtained for different types of stimuli. A survey of the RB literature indicates that RB is almost universally demonstrated with word stimuli, and it is typically very robust in this case (e.g., Bavelier, 1994; Bavelier et al., 1994; Bond & Andrews, 2008; C. L. Harris & Morris, 2000, 2004; Kanwisher & Potter, 1990; Whittlesea et al., 1995; Whittlesea & Masson, 2005). On the other hand, findings of RB for picture stimuli are a lot more variable. RB for pictures of objects has been demonstrated in a number of studies (I. M. Harris & Dux, 2005a, 2005b; I. M. Harris et al., 2012; Hayward et al., 2010; Kanwisher et al., 1999), but the effect size tends to be smaller than for words, and variable across individuals, and is sometimes not significantly different from performance with nonrepeated items. Experiments that manipulated the orientation difference between the critical repeated items found that sometimes there was *less* RB for items that were completely identical (i.e., same object in the same orientation) than when the items differed in orientation and were thus less visually similar, but also more difficult to identify (Hayward et al., 2010). Furthermore, I. M. Harris et al. (2012) found RB for pictures of nonmanipulable objects, but not for pictures of manipulable objects (e.g., tools and other objects associated with specific motor acts); in contrast, word versions of both classes of objects yielded robust RB. This finding suggests that identification of objects from pictures may be influenced by a range of cognitive factors that can modulate susceptibility to RB (see also Goldzieher et al., 2017, for a failure to see RB for pictures of natural scenes). However, from these previous studies, it is difficult to know whether the differences in the presence and strength of the RB phenomenon are due to differences in stimulus format (pictures vs. words), or whether they arise from idiosyncratic differences between the stimulus sets and task requirements of different studies.

The present study had two aims. The first one was to undertake a systematic comparison of RB for word and object stimuli. In order to control for conceptual differences that may have influenced the findings of previous RB studies that used word or object stimuli, we employed word and pictorial

versions of the same concepts, and exactly the same procedure, so that we could directly compare the size of RB across stimulus formats. Our second aim was to revisit the question of whether RB is found when the observers are not required to report the first item of a repeated pair (C1). As reviewed above, there is conflicting evidence for this in previous research using letter and word stimuli. To our knowledge, this question has not been investigated using object stimuli, so, in Experiment 1, we additionally compared report of both the critical items with report of only C2, for both words and pictures of objects.

To anticipate the critical findings, Experiment 1 yielded opposite repetition effects for words versus objects: RB for words and repetition facilitation for objects. Additionally, it found identical patterns of results when participants reported all items and when they reported only the last item of the stream. However, unlike most previous studies of RB, this experiment used a relatively large set of stimuli which were only seen once during the experiment. In order to verify whether this was responsible for the differential repetition effects for words and pictures, Experiment 2 used a more conventional task structure in which a smaller pool of stimuli was presented multiple times during the experiment. This experiment yielded RB for both words and objects, although the effect was substantially stronger for words than for objects. Taken together, these results reinforced the previously observed variability in the RB effect for pictures and are more consistent with an explanation of RB based on difficulties in establishing a robust type representation.

Experiment 1

Method

Participants

Forty psychology undergraduates participated in exchange for partial course credit. They all identified as fluent English speakers/readers, having completed all their schooling in English, and had normal or corrected-to-normal vision. The participants provided informed consent and the procedures were approved by the Human Research Ethics Committee of the University of Sydney. The data from one participant was excluded due to exceedingly low accuracy across all nonrepeat conditions, leaving 39 participants for the analysis.

Apparatus and stimuli

The stimuli consisted of 200 pictures selected from the Hemera Photo-Object database, and their word names. The stimuli comprised a range of object categories spanning both natural kinds and inanimate objects and were chosen to

control a range of word statistics. The 120 critical word items had an average length of 5.7 letters ($SD = 2.07$) and 1.6 syllables ($SD = 0.78$) and an average CELEX frequency of 34.9/million ($SD = 70.41$). The remaining 80 stimuli, which served as intervening filler items, had an average length of 5.8 letters ($SD = 1.90$) and 1.8 syllables ($SD = 0.80$) and an average frequency of a 21.5/million ($SD = 36.96$). The words were presented in Arial font size 50, in black type against a light-gray background. Masks for the word stimuli consisted of strings of symbols (e.g., &#@\$%) in the same font and size. The pictures were converted to grayscale and resized so that their longest dimension was 300 pixels and presented against a light-gray background. They subtended a visual angle of approximately 5° at the viewing distance of ~ 50 cm. Pattern masks were created from fragments of the pictures placed in random orientations and locations which filled a 300×300 pixels square box. The experiment was programmed in Presentation (Neurobehavioral Systems) and was displayed on a 19-in. Dell Trinitron monitor refreshing at 85 Hz.

Design

The experiment comprised two word blocks (full report and partial report) and two picture blocks (full report and partial report), with each type of report using a different list of items, such that each item was only seen once during the experiment in its word or picture form (barring the repetition on repeat trials). The two base lists were carefully constructed to have a spread of items from different word frequency bands and of different syllable and letter lengths across both repeat and nonrepeat trials, ensuring that the lists did not differ on any of these lexical measures. Each participant received the same list for the equivalent block of the opposite format, except with a different randomized trial order, so that the comparisons between words and pictures were always conducted on the same stimuli. Four prerandomized orders were created for each list, yielding 8 different counterbalanced combinations of lists. The order of the word versus picture conditions was counterbalanced across participants, while the order of the partial-report versus full-report conditions was kept fixed within their respective formats, with partial always preceding full, to avoid biasing participants toward actively encoding all items during the partial-report blocks.

Each block consisted of 20 nonrepeat trials, 20 repeat trials, and four catch trials (44 in total, 176 trials across the whole experiment). Each trial started with a 1-s fixation cross, followed by a RSVP stream consisting of three pattern masks, followed by three stimuli (pictures or words, depending on block), followed by three more pattern masks. All stimuli were presented for 106 ms, with no interstimulus interval (see Fig. 1 for examples of trial sequences). The first and third stimuli were considered the critical items (C1 and C2), while the middle item was a filler. C1 and C2 could be the same item

repeated or different items, depending on the condition. Four two-item catch trials in which one of the critical items was replaced by a mask were included in order to discourage participants from assuming the presence of a repetition when they had not detected a third stimulus.

Procedure

Participants were tested individually in a small booth in a testing session lasting approximately 45 minutes. They completed the experiment in two parts, words and pictures, in counterbalanced order. Each of these parts started with a familiarization phase in which the full set of stimuli were presented on-screen one at a time and the participant read or named them out loud. The experimenter corrected any naming errors for the objects and emphasized the name that was considered the correct response and should be used during the experiment. Participants generally made few errors and when they did, these tended to be the name of a very similar item (e.g., “rat” for a picture of a mouse or “turtle” for a picture of a tortoise). Although the correct response was emphasized, such responses were marked as correct during the experimental trial, as they clearly showed the object has been visually identified. This was followed by the two experimental blocks, with the partial-report condition always administered first. For this condition, participants were instructed to name only the last item in the stream. For the full-report block, participants were instructed that two or three items would be presented on each trial, that sometimes a stimulus would be repeated, and they had to name them all, including repetitions. Each of these experimental blocks was preceded by a short practice block (13 trials) containing different stimuli to those used in the experimental block, and with a slightly longer exposure duration of 120 ms per item, to familiarize participants with the RSVP procedure. The experimenter was present in the booth and recorded responses on a laptop. Feedback was given during the practice trials, but not during the experimental blocks.

Results

For the full-report conditions, the measure of interest was the proportion of trials in which both of the critical items were reported correctly (joint C1 and C2 accuracy). This is the most common measure of RB, given that it is not possible to look at individual item accuracy on repeat trials, as one cannot be sure whether participants are reporting C1 or C2 when they only report one. For the partial-report conditions, the measure of interest was the proportion of trials in which the last item was reported correctly. Note that this could inflate C2 accuracy on repeat trials if the participant was in fact reporting C1. Indeed, such a tendency to report C1 or the intervening filler was sometimes observed on nonrepeat trials. Nevertheless, for the purposes of the analysis we assume that the reports reflect

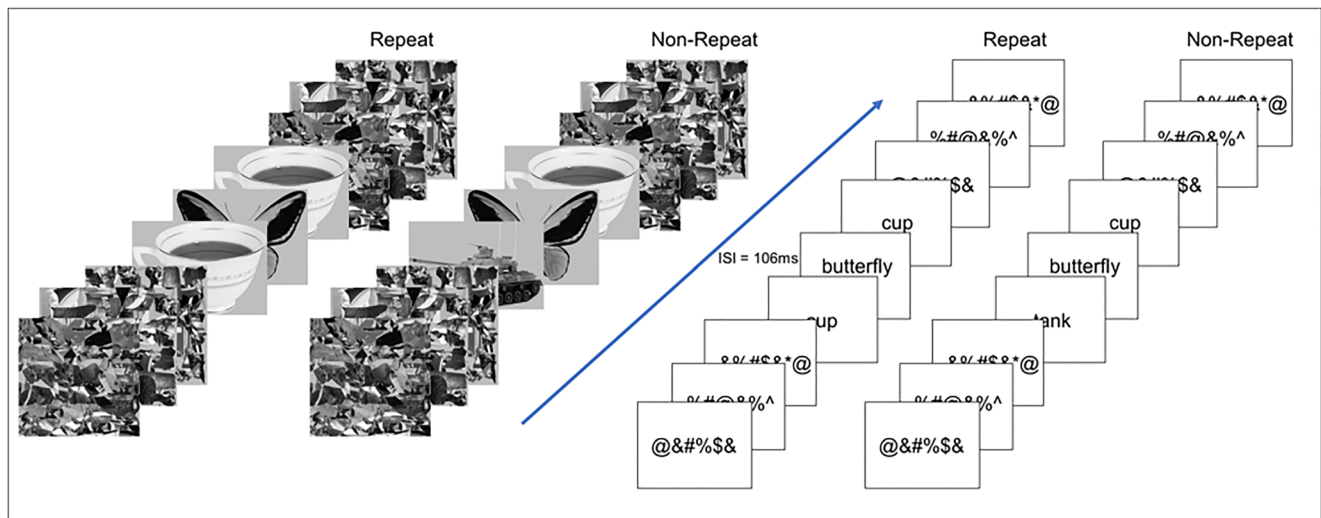


Fig. 1 Examples of the trial structure used in Experiment 1. Three pictures, or words, were preceded and followed by three pattern masks, all presented with an interstimulus interval (ISI) of 106 ms. On half the trials, the first and third items (denoted as C1 and C2) were identical and

on the other half of the trials, they were different items. These were separated by a filler item. A similar structure was used in Experiment 2, except the ISI on different trials could be 94 ms, 106 ms, or 118 ms (see [Method](#) sections for further details)

genuine C2 reports. Figure 2 shows the data for words and pictures, plotted separately for the full-report versus the partial-report conditions.

These data were subjected to a $2 \times 2 \times 2$ repeated-measures analysis of variance (ANOVA), with stimulus format (words vs. pictures), report (full vs. partial), and repetition (repeat vs. nonrepeat) as factors. This showed significantly higher accuracy for words (.64) than for objects (.55) overall, $F(1, 38) = 18.80, p < .001, \eta_p^2 = .331$, but no overall differences between partial-report (.60) and full-report (.58), or between repeat (.60) and nonrepeat (.59) conditions, $F_s < 1, p_s > .64, \eta_p^2 < .005$. Report type did not interact with any of the other variables. However, there was a significant interaction between stimulus format and repetition, $F(1, 38) = 74.09, p < .001, \eta_p^2 = .661$. This was due to the fact that accuracy for words was lower for repeated than for nonrepeated trials, whereas the opposite was true for pictures. Follow-up comparisons

showed significant repetition blindness for words, both under full-report and partial-report conditions ($t_s > 3.51, p_s < .001$, Cohen $d_s > .526$) and significant repetition facilitation for pictures, both under full-report and partial-report conditions ($t_s > 3.88, p_s < .001$, Cohen $d_s > .621$).

In light of the nonsignificant main effect of report in the frequentist ANOVA, a Bayesian analysis was subsequently performed to determine the degree of support for the null hypothesis. This was done using the JASP software (<https://jasp-stats.org>), which uses the Jeffreys–Zellner–Siow (JZS) method to calculate Bayes factors (Rouder et al., 2009). Given inconsistencies in past findings, and the difficulty in obtaining sufficient information from the relevant literature to customize present priors, it is prudent to make minimal assumptions by specifying the mean of the prior distribution as zero (i.e., no difference between partial-report and full-report conditions) and the scale parameter of the Cauchy prior




Fig. 2 Mean proportion of correct reports for words and object pictures in Experiment 1, plotted as a function of repetition condition and type of report. Partial Report = only the last item of the stream (C2) was reported. Full Report = all three items in a stream were reported and the depicted

accuracy is the proportion of trials in which both C1 and C2 were correctly reported. Error bars represent the standard error of the mean difference between the repeat and nonrepeat trials in each block

distribution as 0.707, which are the recommended defaults. For the word condition, the data provide 4.65 times more support for the no-difference model (H_0) over the alternative model (H_1) of there being a difference. For the picture condition, there was 5.69 times more support for H_0 over H_1 . Both of these values qualify as substantial evidence favoring a null effect of report condition (Jarosz & Wiley, 2014).

Discussion

These results show a clear difference in the direction of repetition effects for words versus pictures of objects. Words showed a sizeable RB effect whereas object stimuli yielded significant repetition facilitation. This difference between word and picture stimuli occurred despite the fact that the task used exactly the same stimuli (just in different format) and exactly the same task procedure and response. This suggests that the difference is due to the perceptual attributes of the stimuli and their initial activation of the type, rather than to higher cognitive processes such as memory retrieval or response biases.

Furthermore, reporting only the last item of the stream and reporting all items produced identical patterns of results. The Bayesian analysis conducted on the differences between the two report conditions provides strong evidence favoring the null hypothesis, which gives us confidence that reporting one or both of the critical items does not modulate the repetition effect. Thus, these findings are at odds with Kanwisher's (1987) claim that RB only occurs when there is a need to individuate the two critical items during report and further support the idea that the repetition effects observed here arise from processes involved in the initial activation of types.

The sizeable priming effect found for objects was surprising, given the numerous studies that have demonstrated significant—albeit smaller and more variable—RB with object stimuli. A similar facilitation effect was found by I. M. Harris et al. (2012) in the case of manipulable objects. To check whether this could contribute to the priming found here, we assessed how many objects could be classed as “manipulable” or able to afford a prototypical action. Only 30 out of the 200 could be classed in this way, using a generous criterion of manipulability (anything with a handle [e.g., axe, teapot] or keys [e.g., calculator, piano]), and only between 3–4 out of 20 repeat trials per condition block contained such objects. Thus, it seems very unlikely that this factor could be responsible for the significant priming found.

Another possible explanation is that, because this experiment used a large pool of stimuli, which were only presented once during the experiment, participants had a low degree of familiarity with these items. Morris and Still (2008) argued that more difficult identification conditions, combined with an expectation that stimuli are repeated, encourages guessing

based on partial information gleaned from C1, thus inflating performance on repeat trials.

The competition account put forward by Morris et al. (2009) demonstrated through simulations that when objects are relatively unfamiliar, activating the lexical/name representations of the items can be challenging under RSVP conditions but when objects are presented twice in close temporal proximity, the signal-to-noise ratio of their lexical representations is boosted, leading to facilitation for repeated objects relative to nonrepeated ones. It is, therefore, possible that we observed repetition facilitation rather than RB for objects because our experiment differed from the typical procedure of RB experiments, in which participants are exposed repeatedly to a smaller pool of stimuli and become more familiar with them. This may be particularly important for pictures of objects for which the depiction of the same concept can vary on a wide range of perceptual dimensions.

To test this hypothesis, in Experiment 2, we used a small set of stimuli that were presented several times during the experiment. This experiment also tested three different RSVP presentation rates, to evaluate whether overall task difficulty interacts with differences in repetition effects for words versus pictures.

Experiment 2

Method

Participants

Twenty-four new participants from the same pool took part in this experiment.

Materials

This experiment used a subset of the stimuli from Experiment 1, chosen from amongst the most consistently identified and named items. There were 40 pictures and their corresponding names that served as critical items and 15 pictures and their corresponding names that served as the intervening filler items. The critical word items had an average length of 1.63 syllables ($SD = 0.90$) and an average CELEX frequency of 33.6/million ($SD = 44.50$), and the filler items had an average length of 2.10 syllables ($SD = 0.70$) and an average CELEX frequency of 44.1 ($SD = 77.1$). The masks and display appearance and dimensions were the same as in Experiment 1.

Design

The trial structure was the same as that of Experiment 1: Three stimuli preceded and followed by three masks, except on catch trials, where one of the critical items was replaced by an

additional mask. There were three word conditions and three picture conditions, which differed in presentation rate (94 ms, 106 ms, and 118 ms per stimulus). Each of these conditions consisted of 10 repeat (C1 and C2 were identical), 10 nonrepeat (C1 and C2 were different items) and 10 two-item catch trials (in which C1 or C2 was replaced by a mask, five trials each). For each timing condition, each of the 40 critical items was presented once, with 10 of these being presented twice (on the repeat trials), and each filler item was presented twice. This was replicated across the three timing conditions, using different pairings of critical items and fillers. Across the full complement of word or picture conditions each critical item was seen four times in each format and each intervening filler was seen a total of six times. The repetition conditions and the stimulus presentation rates were all randomly intermixed within separate blocks of words and pictures, respectively, with the order of these blocks counterbalanced across participants.

Procedure

The procedure closely followed that of Experiment 1. Each half of the experiment (words vs. pictures) commenced with a familiarization phase, followed by a short practice block and then the experimental block which consisted of 90 trials. Participants were told that they would see two or three stimuli on each trial, that sometimes a stimulus was repeated and that they had to report all items, including the repetitions. This experiment did not include a partial-report condition.

Results

The proportion of trials in which both C1 and C2 were reported correctly was calculated for each of the conditions and averaged across subjects (see Fig. 3). A $2 \times 2 \times 3$ repeated-measures ANOVA, with the factors stimulus format (words vs. pictures), repetition (repeat vs. nonrepeat) and stimulus duration (94 ms, 106 ms, 118 ms per stimulus) revealed main effects of repetition, $F(1, 23) = 40.59, p < .001, \eta_p^2 = .638$, and

of stimulus duration, $F(1, 22) = 29.13, p < .001, \eta_p^2 = .559$, but no interaction between these factors, $F(1, 22) = 1.75, p = .198$. Although there was no overall difference in accuracy between words and pictures, $F(1, 23) = 1.56, p = .225$, there was a significant interaction between stimulus format and repetition, $F(1, 23) = 18.02, p < .001, \eta_p^2 = .439$. Simple effects analyses showed that words yielded a greater magnitude RB effect ($M = .297, p < .001$) than pictures ($M = .106, p = .005$; see Fig. 3). There was no three-way interaction between repetition, format, and duration, $F(2, 46) = .452, p = .639$.

Comparisons between Experiment 1 and Experiment 2

To check whether the structure of Experiment 2 did, in fact, improve identification accuracy, we compared accuracy in the full-report conditions of Experiment 1 with accuracy in the corresponding 106 ms stimulus duration condition in Experiment 2. In the nonrepeat conditions, accuracy for objects was marginally higher in Experiment 2 than in Experiment 1 (58% vs. 46%), $t(61) = 1.97, p = .053, d = .25$, and a similar trend was apparent for words, although the difference was not significant (78% vs. 68%), $t(61) = 1.66, p = .102, d = .21$. In the repeat conditions, in contrast, accuracy for objects was significantly lower in Experiment 2 than in Experiment 1 (46% vs. 61%), $t(61) = 2.23, p = .030, d = .29$, and accuracy for words was also numerically lower, but again this difference did not reach significance (41% vs. 55%), $t(61) = 1.59, p = .117, d = .20$. The size of the RB effect for words was significantly larger in Experiment 2 compared with Experiment 1 (36 percentage points vs. 13 percentage points), $t(61) = 3.63, p < .001, d = .46$. The same is evidently true for objects, which demonstrated RB in Experiment 2, but repetition facilitation in Experiment 1.

Discussion

Using a smaller pool of stimuli that were presented several times during the experiment, we found RB for both words and




Fig. 3 Mean proportion of correct reports for words and object pictures in Experiment 2, plotted as a function of repetition condition and stimulus duration. Error bars represent the standard error of the mean difference between the repeat and nonrepeat trials in each duration condition

object pictures, although the effect was significantly smaller for objects. These results show that it is possible to elicit RB using the same kinds of stimuli and task requirements as in Experiment 1, and they are consistent with previous findings from experiments with objects that also used smaller pools of stimuli presented multiple times (I. M. Harris & Dux, 2005a, 2005b; I. M. Harris et al., 2012; Hayward et al., 2010; Kanwisher et al., 1999). It is perhaps interesting to note that some of those experiments that found the most robust evidence for RB used even smaller sets of stimuli that were repeated more frequently (I. M. Harris & Dux, 2005a, 2005b), as this suggests that the likelihood of seeing RB for objects is influenced by the number of times a stimulus is encountered during the experiment.

Comparisons between the two experiments provided some evidence (albeit quite weak) that presenting a smaller number of items multiple times made it easier to identify objects on nonrepeat trials. However, there is no convincing evidence that the same improvement was present for words, although it is possible that this reflects a ceiling effect. The perceptual format of common, everyday words is relatively invariant across encounters compared with the variability amongst different images of common objects, so pictures may benefit more from the familiarity gained from repeated exposure in Experiment 2, both in terms of activating their identity representations and the ease of linking those representations with their names. Importantly, though, any increase in familiarity with the stimuli did not translate into an improvement on the repeat trials, for either objects or words, which would be predicted if stimulus identification was generally easier. Instead, performance on repeat trials was substantially lower in this experiment than in Experiment 1. In combination with the increase in accuracy on nonrepeat trials, this translated into a significantly larger RB effect for words in Experiment 2 compared with Experiment 1, as well as a reversal of the repetition facilitation into RB for objects. The latter result may be due to more rapid identification and retrieval of the names of the objects in this experiment, potentially reducing any priming benefit from C1 items that may have obscured an underlying RB effect. The fact that the size of RB was not modulated by exposure duration—even though it did modulate overall accuracy—suggests that the increase in RB is not merely due to a general reduction in task difficulty. Rather, it seems that the difference in patterns of RB between the two experiments is principally due to the use of a smaller number of stimuli that were seen multiple times during the experiment.

General discussion

This study compared the likelihood of eliciting RB for object and word stimuli using identical RSVP tasks that only differed in the format of the stimuli (pictures of objects vs. the word

names of these same objects) and demonstrated a consistent interaction between stimulus format and repetition effects. Across both experiments, words produced robust RB across all conditions tested, including variations in stimulus duration. In contrast, the repetition effects for objects showed considerable fluctuations between experiments. In Experiment 1, which employed a large pool of stimuli each shown once during the experiment (apart from the repetition on a repeat trial), we found significant repetition facilitation for objects; whereas in Experiment 2, which used a smaller subset of stimuli that were each shown 4–6 times during the experiment, we found RB. However, even here the RB effect for objects was significantly smaller in size than the one obtained with words. These findings reinforce many observations in previous research that RB is less consistently observed with object stimuli than with word stimuli and suggest that RB is influenced by the processing dynamics of the type representations.

The results of Experiment 1 also showed identical patterns of performance in the full-report condition, in which participants reported all items in the RSVP stream, and the partial-report condition, in which participants only reported the last item in the stream. A difference between these conditions was used by Kanwisher as one of the main arguments for the token-individuation account, although subsequent studies have failed to replicate that difference (Burt & Jolley, 2017; Kanwisher & Potter, 1990; Leggett et al., 2019; Luo & Caramazza, 1995). Therefore, the strong evidence we found for equivalent patterns of performance between partial-report and full-report conditions is also more in line with an explanation in terms of a failure at the level of type identification than a failure to individuate the items, since the partial-report condition does not enforce token individuation.

A number of factors might be responsible for the observed fluctuations in repetition effects for objects. One possibility is the difficulty of identifying the stimuli in the first place. As outlined in the Introduction, Morris and Still (2008) suggested that difficult task conditions may lead participants to rely on partial information gleaned from poorly processed stimuli and increase guessing, inflating the rate of reporting repeated items and resulting in better performance on repeat trials than nonrepeat trials. This may explain why we found repetition facilitation in Experiment 1, where each object was only seen once, as this could have rendered stimulus identification more difficult. Simulations of Morris et al.'s (2009) competition account showed that processing difficulty induced by short exposure duration was associated with priming rather than RB; similar effects may arise from the difficulty of identifying relatively unfamiliar pictures of objects.

To address this possibility, we tried to make identification easier in Experiment 2 by reducing the number of stimuli and exposing them more often, and by also manipulating the exposure duration. The results partially support Morris and Still's contention, as accuracy on nonrepeat trials with object

stimuli tended to increase under these conditions, and the likelihood of reporting repetitions decreased, resulting in a net RB effect. However, task difficulty is unlikely to be the whole answer, because accuracy for repeat items actually decreased in Experiment 2 relative to Experiment 1, which is not what one would predict if stimulus identification is generally easier. Moreover, there was no conclusive evidence that accuracy for words was affected by the changes in stimulus set composition, although that might be because words are generally very familiar anyway and reading is a more automatic process than object naming. In addition, although overall accuracy improved with longer stimulus duration, the size of the RB in Experiment 2 was not modulated by exposure duration. It is possible, however, that the increased familiarity with the items increased the competition between items with strong representations and increased the chance of seeing an underlying RB effect (Morris et al., 2009).

Another possible contributor to the different repetition effects for objects is that Experiment 2 contained a higher proportion of two-item trials (33% of all trials) than Experiment 1 did (10%). It is conceivable that experiencing many trials in which a third item was genuinely missing led participants to adopt a more conservative criterion for reporting repetitions. The results for the words could be consistent with a shift in response criterion, as the size of the RB effect for words was larger in Experiment 2 (36 percentage points for the 106-ms condition, 29 percentage points for the 94 ms, and 24 percentage points for the 118-ms conditions) than in Experiment 1 (13 percentage points for the corresponding 106-ms full-report condition). This may also account for why we observed RB for objects in Experiment 2, but not in Experiment 1. However, this explanation does not sit well with the fact that all the previous experiments that found RB for objects used only a small number of two-item trials (generally 10% of all trials), similar to Experiment 1 (I. M. Harris & Dux, 2005a, 2005b; I. M. Harris et al., 2012; Hayward et al., 2010; Kanwisher et al., 1999). Clearly then, while the number of catch trials may contribute to the difference in RB effects for objects between the two experiments, it is not the whole explanation.

The final possibility is that the number of times participants were exposed to the individual stimuli during the experiment influenced the pattern of repetition effects for object pictures across the two experiments. How would this factor affect RB? There is good evidence that initial identification is accomplished on the basis of component features (e.g., object parts), which are activated early during visual processing (Grill-Spector & Kanwisher, 2005; Ullman et al., 2002); such features are often sufficient to uniquely identify objects, at least at a categorical level, and are thought to be the main way in which objects are recognized during RSVP (I. M. Harris et al., 2010; I. M. Harris et al., 2008; Hayward et al., 2010). However, identification on the basis of component features

can be challenging when the same features are activated repeatedly, because there is insufficient time to establish that these features belong to separate object-level representations (object tokens), leading to RB (Hayward et al., 2010). This problem could be exacerbated in an experimental context where the same items are presented multiple times, potentially creating a situation where the participant is unsure whether they had seen a stimulus on the current trial or a previous one. Essentially, according to this view, experimental contexts with multiple presentations of the same items throughout the experiment are likely to induce greater RB, because of greater ambiguity both at a featural level—which prevents the formation of configural representations (i.e., object-level representations) that are necessary to disambiguate the feature-level representations and uniquely identify objects (Kent et al., 2016)—and at the object level.

Note that this explanation also applies to the results with word stimuli. For words, there is *always* high featural ambiguity and featural noise, because words share the same small set of 26 letters, which themselves are made up of an even smaller set of low-level features. This idea is consistent with previous findings that RB scales with the amount of orthographic overlap (shared letters and letter clusters) between the critical words (C. L. Harris & Morris, 2000), and could explain why words almost universally show reliable and robust RB. In the current study, the size of RB for words was higher in Experiment 2 than in Experiment 1, in line with an increase in the ambiguity at the word level. Therefore, while we cannot discount the influence of task difficulty or shifts in response criterion on the present results, we propose that ambiguity at different levels of the representational hierarchy (from feature to object levels) plays as large, if not a larger, role in RB.

This account of RB is admittedly post hoc and needs to be tested in future research. However, it has the added advantage of also providing an explanation for some findings that are harder to reconcile with an account based purely on token individuation. Namely, a range of findings show that an RB-like deficit occurs even for items that are only similar, rather than identical. This includes orthographically similar words that share most of their letters (e.g., *must* and *gust*; *cap* and *cape*; Bavelier et al., 1994; C. L. Harris & Morris, 2000; Kanwisher & Potter, 1990), visually similar pictures of objects (e.g., pear and guitar; Kanwisher et al., 1999), and objects that overlap in visuosemantic features (e.g., horse and camel; Seet et al., 2019). These stimuli do not share the same type, so an explanation in terms of a failure of token individuation of the same activated type is problematic. These findings had variously been explained in terms of activation of sublexical units, or a shared “semantic category” type, all of which stretch the notion of “type” in rather arbitrary ways.

Instead, by acknowledging that these stimuli have a high degree of shared features, our account provides a unifying explanation for these findings that relies on the same

explanation as the basic finding of RB for identical repetitions. The fundamental problem might really lie in a failure to form a full and stable representation of the type, particularly as it pertains to the object level of representation that is tied to conscious awareness, and the experimental context may be a factor that contributes to the difficulty of forming these representations.

In summary, the present study has revealed substantial differences between repetition blindness for words and object stimuli, and found equivalent patterns of repetition across partial-report and full-report RSVP tasks. Taken together, these findings emphasize the need to rethink the token-individuation framework and suggest instead that RB arises due to difficulties in establishing a robust type representation.

Acknowledgements This research was supported by Australian Research Council grant DP0879206.

Open practices statement The data and materials for all experiments reported here are available from the authors on request.

References

- Bavelier, D. (1994). Repetition blindness between visually different items: The case of pictures and words. *Cognition*, 51(3), 199–236. [https://doi.org/10.1016/0010-0277\(94\)90054-x](https://doi.org/10.1016/0010-0277(94)90054-x)
- Bavelier, D., Prasada, S., & Segui, J. (1994). Repetition blindness between words: Nature of the orthographic and phonological representations involved. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(6), 1437–1455. <https://doi.org/10.1037/0278-7393.20.6.1437>
- Bond, R., & Andrews, S. (2008). Repetition blindness in sentence contexts: Not just an attribution? *Memory & Cognition*, 36, 295–313. <https://doi.org/10.3758/MC.36.2.295>
- Burt, J. S., & Jolley, J. (2017). Repetition blindness in priming in perceptual identification: Competitive effects of a word intervening between prime and target. *Memory & Cognition*, 45(7), 1171–1181. <https://doi.org/10.3758/s13421-017-0726-z>
- Fagot, C., & Pashler, H. (1995). Repetition blindness: Perception or memory failure? *Journal of Experimental Psychology: Human Perception and Performance*, 21, 275–292. <https://doi.org/10.1037/0096-1523.21.2.275>
- Goldzieher, M. J., Andrews, S., & Harris, I. M. (2017). Two scenes or not two scenes: The effects of stimulus repetition and view-similarity on scene categorization from brief displays. *Memory & Cognition*, 45, 49–62. <https://doi.org/10.3758/s13421-016-0640-9>
- Grill-Spector, K., & Kanwisher, N. (2005). Visual recognition: As soon as you know it is there, you know what it is. *Psychological Science*, 16, 152–160. <https://doi.org/10.1111/j.0956-7976.2005.00796.x>
- Harris, C. L., & Morris, A. L. (2000). Orthographic repetition blindness. *Quarterly Journal of Experimental Psychology Section A*, 53A(4), 1039–1060. <https://doi.org/10.1080/02724980050156281>
- Harris, C. L., & Morris, A. L. (2004). Repetition blindness occurs in nonwords. *Journal of Experimental Psychology: Human Perception and Performance*, 30, 305–318. <https://doi.org/10.1037/0096-1523.30.2.305>
- Harris, I. M., Benito, C. T., & Dux, P. E. (2010). Priming from distractors in rapid serial visual presentation is modulated by image properties and attention. *Journal of Experimental Psychology: Human Perception and Performance*, 36(6), 1595–1608. <https://doi.org/10.1037/a0019218>
- Harris, I. M., & Dux, P. E. (2005a). Orientation-invariant object recognition: Evidence from repetition blindness. *Cognition*, 95, 73–93.
- Harris, I. M., & Dux, P. E. (2005b). Turning objects on their head: The influence of the stored axis on object individuation. *Perception and Psychophysics*, 67(6), 1010–1015. <https://doi.org/10.3758/BF03193627>
- Harris, I. M., Dux, P. E., Benito, C. T., & Leek, E. C. (2008). Orientation sensitivity at different stages of object processing: Evidence from repetition priming and naming. *PLOS ONE*, 3(5), Article e2256. <https://doi.org/10.1371/journal.pone.0002256>
- Harris, I. M., Murray, A. M., Hayward, W. G., O’Callaghan, C., & Andrews, S. (2012). Repetition blindness reveals differences between the representations of manipulable and nonmanipulable objects. *Journal of Experimental Psychology: Human Perception and Performance*, 38, 1228–1241. <https://doi.org/10.1037/a0029035>
- Hayward, W. G., Zhou, G., Man, W.-F., & Harris, I. M. (2010). Repetition blindness for rotated objects. *Journal of Experimental Psychology: Human Perception and Performance*, 36(1), 57–73. <https://doi.org/10.1037/a0017447>
- Jarosz, A. F., & Wiley, J. (2014). What are the odds? A practical guide to computing and reporting Bayes factors. *The Journal of Problem Solving*, 7(1), Article 2. <https://doi.org/10.7771/1932-6246.1167>
- Kanwisher, N. (1987). Repetition blindness: type recognition without token individuation. *Cognition*, 27(2), 117–143. [https://doi.org/10.1016/0010-0277\(87\)90016-3](https://doi.org/10.1016/0010-0277(87)90016-3)
- Kanwisher, N., & Potter, M. C. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance*, 16(1), 30–47. <https://doi.org/10.1037/0096-1523.16.1.30>
- Kanwisher, N., Yin, C., & Wojciulik, E. (1999). Repetition blindness for pictures: Evidence for the rapid computation of abstract visual descriptions. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp. 119–150). Cambridge, Mass.: MIT Press.
- Kent, B. A., Hvoslef-Eide, M., Saksida, L. M., & Bussey, T. J. (2016). The representational–hierarchical view of pattern separation: Not just hippocampus, not just space, not just memory? *Neurobiology of Learning and Memory*, 129, 99–106. <https://doi.org/10.1016/j.nlm.2016.01.006>
- Leggett, J. M. I., Burt, J. S., & Ceccato, J.-M. (2019). Repetition priming and repetition blindness: Effects of an intervening distractor word. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 73(2), 105–117. <https://doi.org/10.1037/cep0000164>
- Luo, C. R., & Caramazza, A. (1995). Repetition blindness under minimum memory load: Effects of spatial and temporal proximity and the encoding effectiveness of the first item. *Perception and Psychophysics*, 57, 1053–1064. <https://doi.org/10.3758/BF03205464>
- 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(1), 95–113.
- Morris, A. L., & Still, M. L. (2008). Now you see it, now you don’t: Repetition blindness for nonwords. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 146–166. <https://doi.org/10.1037/0278-7393.34.1.146>
- Morris, A. L., Still, M. L., & Caldwell-Harris, C. L. (2009). Repetition blindness: An emergent property of inter-item competition. *Cognitive Psychology*, 58, 338–375. <https://doi.org/10.1016/j.cogpsych.2008.08.001>
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2), 225–237. <https://doi.org/10.3758/PBR.16.2.225>

- Seet, M. S., Andrews, S., & Harris, I. M. (2019). Semantic repetition blindness and associative facilitation in the identification of stimuli in rapid serial visual presentation. *Memory & Cognition*, 47(5), 1024–1030. <https://doi.org/10.3758/s13421-019-00905-9>
- Ullman, S., Vidal-Naquet, M., & Sali, E. (2002). Visual features of intermediate complexity and their use in classification. *Nature Neuroscience*, 5, 682–687. <https://doi.org/10.1038/nn870>
- 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, & Cognition*, 21(6), 1670–1688. <https://doi.org/10.1037/0278-7393.21.6.1670>
- Whittlesea, B. W. A., & Masson, M. E. J. (2005). Repetition blindness in rapid lists: Activation and inhibition versus construction and attribution. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 54–67. <https://doi.org/10.1037/0278-7393.31.1.54>
- Whittlesea, B. W. A., & Podrouzek, K. W. (1995). Repeated events in rapid lists: Part 2. Remembering repetitions. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 21, 1689–1697. <https://doi.org/10.1037/0278-7393.21.6.1689>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.