

Repetition Blindness: Levels of Processing Revisited

Alison L. Morris (alisonm@bu.edu)
Department of Psychology, Boston University
64 Cummington St., Boston, MA 02215

Catherine L. Harris (charris@bu.edu)
Department of Psychology, Boston University
64 Cummington St., Boston, MA 02215

Abstract

When two orthographically similar words are briefly and successively displayed, the second word is often difficult to detect or recall, a deficit known as repetition blindness, or RB (Kanwisher, 1987). Two experiments used word-nonword pairs to test predictions of a computational model based on similarity inhibition (Bavelier & Jordan, 1992) vs. predictions of a sublexical model (Harris & Morris, 1996, 1997; Morris & Harris, 1997). One striking finding was of strong RB even for a single repeated letter (*cope carn; hot hix*). Results generally supported a sublexical model where only the shared letters are affected by RB, and each shared letter can be differentially affected in a probabilistic manner.

Introduction

Repetition Blindness (RB), the failure to detect repetitions of visual events in Rapid Serial Visual Presentation, is a striking and robust phenomenon which has thus far eluded complete explanation despite extensive study. Early investigations of RB using RSVP sentence stimuli (Kanwisher, 1987; Kanwisher & Potter, 1990) demonstrated that, when shown a sentence such as *Her jacket was red because red is conspicuous*, subjects would typically report, *Her jacket was red because is conspicuous*. RB is time- and exposure-dependent; the effect is eliminated with increased exposure of the repeated word, or increased time between repetitions (Kanwisher, 1987; Park & Kanwisher, 1994). RB has been found for repeated words or pictures, orthographically similar words, and phonologically similar items such as *nine* and *9* (Bavelier & Potter, 1992; Bavelier, 1994). The phenomenon of RB impacts on models of high-level visual processing, word recognition, and consciousness, and therefore is of considerable relevance to cognitive scientists.

One prevalent explanation of RB emphasizes the distinction between types and tokens in visual processing (Kanwisher, 1987; Kanwisher & Potter, 1990; Kanwisher, 1991; Park & Kanwisher, 1994). In this view, referred to as "type activation without token individuation", a word's recognition node (type) could be activated twice, but only one episodic representation (token) is formed. Although this explanation has intuitive appeal, a formal, computational account of the types-and-tokens model has yet to emerge. In addition, it is difficult to see how this model accounts for RB between non-identical words. Kanwisher and Potter (1990) suggested that RB between orthographically similar

words comes about because one word is misread as the other. These investigators have alternatively suggested, along with Bavelier, Prasada, and Segui (1994) that RB effects might be located at the level of abstract letter clusters; however, this idea has not been proposed as a formal part of the types-and-tokens model.

Other investigators have attempted to locate the mechanism responsible for RB within the word recognition system (Bavelier & Jordan, 1992; Bavelier, Prasada, & Segui, 1994). Bavelier and Jordan proposed a computational model of visual word recognition that accounts for RB as well as other time-dependent phenomena such as masked priming. In their model, detection of a word depends on levels of activation relative to a variable baseline and threshold. When a word is displayed for sufficient time for it to pass its threshold and be detected, its baseline is then temporarily raised. RB results when a second display of the word occurs while its baseline is still elevated from the first detection, because in this case, the level of activation relative to baseline will be treated as noise. The model attempts to account for orthographic RB by also raising the baselines of all other words in proportion to their similarity to the detected word.

Our own investigations of RB phenomena (Harris & Morris, 1996; 1997; Morris & Harris, 1997) have suggested that RB between non-identical words is not a matter of similarity inhibition, as in the Bavelier and Jordan model, but rather, that when RB occurs between orthographically similar words (called W1 and W2), only the shared letters of W2 are affected; the unique letters are unaffected and remain available for activating words. The differential effect on shared vs. nonshared letters has been most strikingly demonstrated using an illusory words paradigm (Harris & Morris, 1996; Morris & Harris, 1997). In this paradigm, RSVP streams of words and word fragments are constructed such that the W2's unique letters can combine with a subsequent word fragment to form an "illusory" word. Thus, display of *light rock shock ell cup* results in report of *light rock shell cup*; further, reports of illusory words are more prevalent in such "repeated" conditions compared to unrepeated conditions designed to produce illusory words through simple letter migration. Other evidence for the sublexical nature of RB between orthographically similar words comes from trials where subjects misreport W2 as some other word. These misreports are statistically more

likely to preserve the word's unique letters, and omit the repeated letters, than in control conditions where W1 and W2 share no letters (Harris & Morris, 1997).

The work cited above is more consistent with sublexical explanations of RB than with word-to-word inhibition accounts such as that of Bavelier and Jordan (1992). However, details of the sublexical account remain to be elucidated. One important question is the nature of the sublexical units. We will discuss two sublexical accounts: the single letter and the letter clusters accounts.

Sublexical units are individual letters

The simplest account is probably that the sublexical units are individual letters. However, failure to find RB for anagrams such as *early* and *layer* (Kanwisher, 1986) indicates that the individual letters need to have some type of position-specific or contextual coding. Kanwisher and Potter (1990) tested the hypothesis that display of stimuli sharing one letter in the same position, such as *fault* and *heart*, would lead to "stripping" of the shared letter from W2, resulting in report of *fault hear*. Their failure to find such letter stripping led them to conclude that RB between similar words is not the result of independent letter-level effects.

It is possible that RB does stem from letter-level effects, but that this was obscured by Kanwisher and Potter's choice of stimulus materials. They used critical words which shared only the last letter. Subjects might have been blind to the identity of the repeated letter, but retained some notion of word length. The target word could have been correctly generated via pattern completion, as happens when letters in a word are "missing" because they are covered by an ink blot (McClelland & Rumelhart, 1981). We reasoned that pattern completion might be a less favored solution if the repeated letter were at the beginning of the word (Harris & Morris, 1997). Using materials like *bluff BRACE*, *clock CHARM*, and *doodle DANGER*, we found two types of evidence for repetition blindness to a single letter. Subjects would sometimes say "*doodle anger*", but also frequently reported a similar-length word in which the repeated letter had been replaced with another letter, such as reporting *bluff grace* instead of *bluff brace*. For this reason, the status of the independent letter-level explanation for RB remains unresolved.

Letter clusters account

A more complicated model is that the relevant units are context-dependent, overlapping letter clusters such as the digrams and trigrams *dr*, *ru*, *rug*, *ugs* found in *drugstore* (Mozer, 1991; Harris & Morris, 1996; 1997). The letter clusters account also explains the failure of Kanwisher and Potter to find letter-stripping for stimuli like *fault* and *heart*. The sublexical units in *heart* include the digram units *rt* and *t_* (where "_" indicates adjacency to a word boundary). Only *t_* is shared by both words. The presence of *rt* means that the final letter *t* in *heart* has some probability of being reconstructed from *rt*. Key evidence that letter clusters are the units which mediate visual word recognition comes from studies on letter migration (Mozer, 1983). However, letter cluster representations are generally useful for ex-

plaining readers' high degree of sensitivity to letter co-occurrence frequencies (Seidenberg & McClelland, 1989). That is, words containing highly frequent clusters, such as *fake* and *crock*, are easy to read, even when the words themselves are low frequency. The ease with which humans read pronounceable nonwords may also result from having unitized letter cluster representations (Plaut, McClelland, Seidenberg, & Patterson, 1996).

Empirically distinguishing between the word-level and sublexical accounts is made difficult by the fact that word fragments activate word representations (McClelland & Rumelhart, 1981). The sublexical hypothesis proposes that only the repeated units are affected by RB. If RB is assessed by ability to report W2, the amount of RB will indirectly depend on whether the unaffected units strongly or weakly evoke the target word. For this reason, nonwords may provide a "purer" medium for assessing RB, as they are less affected by such pattern completion mechanisms.

The experiments described in this paper used successively displayed word-nonword pairs of varying degrees of orthographic similarity, with the word as W1 and the nonword as W2. Experiment 1 was designed to test a prediction of similarity inhibition models against a prediction of the most simple form of sublexical model.

Experiment 1

Similarity inhibition models predict that a word and nonword sharing all but one of their letters (*card carn*) would show substantially greater RB than a word and nonword sharing only one letter (*cope carn*). In contrast, a sublexical model assuming simple refractoriness or token individuation failure at the (position-specific) individual letter level predicts that veridical report of W2 will be equally difficult for *cope carn* as for *card carn*. In both cases, suppression of the repeated letters will leave a fragment which will either be reported or will be subjected to pattern completion, causing subjects to report something other than the correct stimulus.

Materials and Procedure

Subjects were 15 Boston University students who participated in exchange for course credit. All subjects were native English speakers (3 subjects acquired English simultaneously with another language).

Twenty-one easily pronounceable 3- or 4-letter nonwords were created and each was paired with three different words to create three versions of each stimulus item. In the **Neighbor** condition, the word and nonword shared all but their final letters (the first two letters were shared in the case of 3-letter word-nonword pairs, with the first three letters shared in the case of 4-letter word-nonword pairs). In the **Letter** condition, only the initial letters were shared; and in the **Unrepeated** condition, no letters were shared. The words in each stimulus item were matched for frequency across conditions. Examples of the three stimulus conditions are shown in Table 1 below. The three versions of each stimulus item were counterbalanced across subjects, such that each subject viewed 7 stimulus items in each condition, for a total of 21 experimental trials. The word

(W1) always appeared in lower case letters, with the nonword (W2) in upper case. Stimuli were arranged in 7 blocks of 5 trials each; the blocks each contained 1 Neighbor, 1 Letter, and 1 Unrepeated word-nonword pair plus two filler trials containing no repeated letters.

Table 1: Examples of Stimuli (Experiment 1)

Condition	W1	W2
Neighbor	mud	MUP
	gate	GATH
Letter	mad	MUP
	golf	GATH
Unrepeated	row	MUP
	knee	GATH

Subjects were instructed to report both the word and nonword on each trial. Pilot studies indicated that subjects varied considerably in the ability to perceive and report briefly displayed nonwords; subjects were therefore instructed to report whatever letters they did perceive in the event that they were unable to report an entire nonword. A lengthy practice procedure was also instituted in order to improve nonword report. Subjects viewed 30 practice trials consisting of word-nonword pairs (with no repeated letters) prior to viewing the experimental trials. Each trial began with a "+" appearing in the center of the computer screen as a fixation point. When the subject pressed the space bar, the "+" disappeared and the word appeared 500 msec later. The word was displayed for 90 msec, followed immediately by the nonword (there was an approximately 15 msec delay in the onset of the nonword, due to the refresh rate of the computer monitor). Exposure duration for the nonword was initially set at 195 msec, and was subsequently decreased by 15 msec every 5 trials if the subject's accuracy in reporting both the word and nonword was greater than 60% over the preceding block of 5 trials. Exposure duration was increased if the subject's accuracy fell below 45%. The average exposure duration for the nonword for the experimental trials across the 15 subjects was 108 msec., range 75-195 msec. The nonword was immediately followed by a mask (row of ampersands) displayed for 250 msec. All stimuli were centered on the display. Experimenters recorded via keypresses whether subjects reported both critical items, W1 only, W2 only, or neither of the critical items. Experimenters also annotated the scoresheets in the event that the subject reported a different word or nonword than the ones displayed, or reported only part of the word or nonword. The stimuli were presented on a Macintosh IICI, controlled by PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993). The font was 36 pt. Courier. Subjects sat 20 inches from the screen.

Results and Discussion

Percent correct report of W1 averaged 97% or greater in all three stimulus conditions. Table 2 shows percent correct report of W2 given correct report of W1. Analysis of vari-

ance revealed a highly significant effect of condition in both the subject and items analyses, $F_1(2, 28) = 44.4, p < .001, \eta^2 = .76$; $F_2(2, 40) = 51.2, p < .001, \eta^2 = .72$.

Table 2: Percent Correct Report of W2 given Report of W1 (Experiment 1)

Condition	mean	S.E.
Neighbor	.11	.03
Letter	.43	.06
Unrepeated	.72	.05

Newman-Keuls tests showed that report of W2 was significantly lower for both Neighbor and Letter conditions compared to the Unrepeated condition (both $ps < .01$). Since RB is indexed by the difference between repeated (Neighbor or Letter) and Unrepeated conditions, these results showed a substantial repetition deficit for both Neighbor and Letter conditions. In addition, report of W2 in the Neighbor condition was lower than in the Letter condition, ($p < .01$) indicating greater RB in the Neighbor condition (61%) compared to the Letter condition (29%).

Models of RB based on similarity inhibition predict that word-nonword pairs sharing all but one letter (Neighbor condition) would show greater RB than word-nonword pairs sharing only the first letter (Letter condition). The results of Experiment 1 are therefore consistent with a similarity inhibition model of RB. On the other hand, closer examination of subjects' error reports reveals evidence inconsistent with word-to-word inhibition, but consistent with sublexical models.

Table 3 shows the percentage of misreports of W2 relative to correct reports and omissions. Misreports include reporting a word, another nonword, or individual letters (as in, for the stimulus *box BOT*, reporting "*box and something ending in T*"). Such misreports formed a substantial portion of W2 error reports in the two repeated conditions.

Table 3: Percent Correct, Omission, and Misreport of W2 (Experiment 1)

Condition	Correct	Omission	Misreport
Neighbor	.11	.63	.26
Letter	.43	.27	.30
Unrepeated	.72	.14	.14

While word-to-word inhibition assumes that W2 will be frequently omitted from report in repeated conditions, this type of model makes no predictions about misreporting of W2, since the entire word is assumed to be affected by RB. In contrast, sublexical models predict the occurrence of a particular type of misreport: Letters *unique* to W2 will be reported (either as part of another word, or in isolation) while letters *shared* with W1 will be excluded.

Specifically, sublexical models predict that, for the Letter condition, the first letter of W2 (which is shared with W1) is likely to be excluded from the misreport, while the remaining letters are likely to be preserved, relative to the Unre-

peated condition. The top half of Table 4 shows percent report of the first letter vs. the remaining letters for the Letter and Unrepeated conditions. As predicted by sublexical models, W2 misreports in the Letter condition selectively omitted the letter shared with W1, as compared to the same letter in the Unrepeated condition, $t(23) = 4.0, p < .001, \eta^2 = .41$. Similarly, the unique letters in W2 were reported at a higher rate in the Letter condition compared to the same letters in the Unrepeated condition, $t(23) = 4.8, p < .001, \eta^2 = .50$.

Table 4: Percent Report of Shared vs. Unique Letters in W2 (Experiment 1)

Condition	%First letter	%Remaining letters	Examples	
			Stimulus	Misreport
Letter	.14	.94	win wep	win hep
Unrepeated	.72	.56	bay wep	bay weep
Condition	%shared letters	%last letter	Examples	
			Stimulus	Misreport
Neighbor	.44	1.00	noon noof	noon half
Unrepeated	.65	.52	drug noof	drug no

The bottom half of Table 4 shows a similar comparison involving the Neighbor and Unrepeated conditions. In this case, the Neighbor condition has only one unique letter, the last letter. All other letters are shared. Corresponding letters in the Unrepeated condition would be the first two letters in the case of 3-letter nonwords, and first 3 letters in the case of 4-letter nonwords. Again, the shared letters appeared less often in misreports in the Neighbor condition when compared with the same letters in the Unrepeated condition, $t(25) = 2.4, p < .05, \eta^2 = .19$, while the unique letter was selectively preserved in the Neighbor condition in comparison with the same letter in the Unrepeated condition, $t(8) = 2.9, p < .02, \eta^2 = .51$. We can also examine the form of the misreport errors in more detail. Table 5 shows percent report of four error types as a percentage of all W2 misreport errors for each condition: reporting a word in place of the nonword, or word conversion (*meat MEAP* ---> *meat CARP*); report of another nonword of approximately the same length (*joy JEG* ---> *joy ZEG*); letter stripping (*hot HIX* ---> *hot IX*); and report of other isolated letters, not letter stripping (*date RISP* ---> *date "R"*).

Table 5: Misreport Types (Experiment 1)

Condition	Word Conversion	Other Nonword	Letter Stripping	Other Letters
	Neighbor	.47	.47	0
Letter	.39	.24	.32	.05
Unrepeated	.52	.39	0	.09

Most misreports were either word conversions or reports of another nonword. Although letter stripping misreports

did occur fairly frequently in the Letter condition, it should be noted that 57% of these letter strippings also could be classified as word conversions (such as *mad MUP* ---> *mad UP* and *bird BEEL* ---> *bird EEL*). The error analysis from Experiment 1 is therefore consistent with the results of our previous analysis using word stimuli (Harris & Morris, 1997) in terms of both the selectivity of RB (affecting shared but not nonshared letters) and misreport type (letter stripping vs. word-level pattern completion).

The overall results of Experiment 1 would seem to be somewhat ambiguous, as the error analysis points to a sublexical locus, but the finding of more RB for the Neighbor condition compared to the Letter condition is more consistent with a similarity inhibition model. However, this experiment contained a confound between the *proportion* of shared letters between the Neighbor and Letter conditions, and the *number* of shared letters. This confound may explain the ambiguous results. The simplest sublexical model assumes that when W2 is affected by RB, all shared letters in that W2 are equally affected. On the other hand, if we assume that each repeated letter has a separate probability of being suppressed by RB, then we would indeed expect more RB for the Neighbor condition than for the Letter condition, because with each additional shared letter, the probability of all letters being correctly perceived decreases. In Experiment 2, we re-examined the question of word-to-word inhibition vs. sublexical locus using stimuli which shared the same number, but different proportions, of their letters across conditions.

Experiment 2

Experiment 2 again used word-nonword pairs to investigate the locus of the RB effect. In this experiment, we held the number of shared letters constant across items (in the Repeated condition) but varied the proportion of shared letters by varying word length. This enabled us to manipulate similarity between W1 and W2 and resulted in different predictions for similarity inhibition vs. sublexical models of RB. Words and nonwords such as *jet* and *JEG* are more similar to each other than *plane* and *PLOSH*, although both share the same number of letters. Similarity inhibition accounts therefore predict a greater magnitude of RB for *jet JEG* compared to *plane PLOSH*, while sublexical accounts predict similar amounts of RB.

A methodological change was also introduced in Experiment 2. In Experiment 1, it is possible that, with no lag separating W1 and W2 (other than the 15-30 msec. refresh time) our results could have been influenced by visual fusion, even though W1 and W2 were displayed in different case. In order to eliminate this possibility, a row of symbols was displayed for 60 msec. between W1 and W2, which introduced a lag of approximately. 75-90 msec.

Materials and Procedure

Subjects were 20 Boston University students who participated in exchange for course credit. All subjects were native English speakers.

Stimuli for Experiment 2 consisted of 20 word-nonword pairs, with two different lengths: **three** letters and **five** let-

ters. Within each stimulus type there was also a Repeated and Unrepeated condition. For the **Repeated** condition, the word and nonword shared the first two letters; no letters were shared in the **Unrepeated** condition. Thus, all **Repeated** items shared the same number of letters between W1 and W2, but the proportion of letters shared was greater for the 3-letter stimuli (67%) than for the 5-letter stimuli (40%). W1's were again matched for frequency between conditions and stimulus types. Examples of stimuli for Experiment 2 are shown in Table 6.

Table 6: Examples of Stimuli (Experiment 2)

Condition	W1	W2	Condition	W1	W2
3-letters			5-letters		
Repeated	cup	CUG	Repeated	phone	PHURL
Unrepeated	raw	CUG	Unrepeated	stone	PHURL

Stimuli were arranged in 5 blocks of 6 trials each. The blocks each contained 1 Repeated and 1 Unrepeated 3-letter word-nonword pair, 1 Repeated and 1 Unrepeated 5-letter word-nonword pair, and two filler trials with no repeated letters. Thus, each subject viewed 5 stimulus items of each of the four types, plus 10 filler trials. Stimuli were rotated through the various blocks, creating 10 different presentation orders. The word (W1) always appeared in lower case letters, with the nonword (W2) in upper case.

Subjects again viewed 30 practice trials prior to the experimental trials. Procedure was the same as in Experiment 1, except that a row of pound signs (#####) was displayed for 60 msec following W1 and prior to display of W2. Exposure duration for the nonword (W2) was set during the practice trials as described in Experiment 1; duration was held constant following the practice trials. The average exposure duration for W2 for the experimental trials across the 20 subjects was 95 msec., range 75-135 msec.

Results and Discussion

Correct report of W1 averaged 98% or greater in all stimulus conditions. Percent correct report of W2 given correct report of W1 is shown in Table 7. Analysis of variance revealed a significant main effect of repeatedness in both the subject and item analyses, $F_1(1, 19) = 26.6, p < .001, \eta^2 = .44$; $F_2(1, 18) = 67.1, p < .001, \eta^2 = .50$. The main effect of length was not significant, $F_1(1, 19) = 2.3, p > .10, \eta^2 = .01$; $F_2(1, 18) = 1.5, p > .20, \eta^2 = .02$. We also failed to obtain a significant interaction between repeatedness and length, $F_1(1, 19) = 1.2, p > .25, \eta^2 = .01$; $F_2(1, 18) = 1.1, p > .30, \eta^2 = .01$.

Table 7: Percent Correct Report of W2 given Report of W1 (Experiment 2)

Condition	mean	S.E.	Condition	mean	S.E.
3-letters			5-letters		
Repeated	.31	.07	Repeated	.29	.06
Unrepeated	.72	.06	Unrepeated	.61	.06

Table 8 shows the percentage of misreports of W2 relative to correct reports and omissions, collapsed across stimulus length. Misreports include reporting a word, another nonword, or individual letters.

Table 8: Percent Correct, Omission, and Misreport of W2 (Experiment 2)

Condition	Correct	Omission	Misreport
Repeated	.30	.24	.46
Unrepeated	.67	.09	.24

As in Experiment 1, sublexical models predict that in misreports, letters unique to W2 will be reported, while letters shared with W1 will be excluded from report. In Experiment 2, the first two letters in the Repeated condition are shared; the remaining letters are unique to W2. Table 9 shows the percent report of the first 2 letters vs. the remaining letters for the Repeated and Unrepeated conditions, collapsed across word lengths. As predicted by sublexical models, W2 misreports in the Repeated condition selectively omitted the letters shared with W1 as compared to the same letters in the Unrepeated condition, $t(17) = 6.2, p < .001, \eta^2 = .69$. Also, the unique letters in W2 were reported at a higher rate in the Repeated condition compared to the same letters in the Unrepeated condition, $t(17) = 2.8, p < .02, \eta^2 = .32$.

Table 9: Percent Report of Shared vs. Unique Letters in W2 (Experiment 2)

Condition	%First 2 letters	%Unique letters	Examples	
			Stimulus	Misreport
Repeated	.36	.80	sweet swand	sweet rand
			mud mup	mud wup
Unrepeated	.69	.54	thick swand	thick swamp
			row mup	row mut

Table 10 shows percent report of word conversions, other nonword report, letter stripping, and report of other letters (non-letter stripping) as a percentage of all W2 misreport errors for each condition. This analysis shows similar results to Experiment 1: some letter stripping did occur along with the other misreport types.

Table 10: Misreport Types (Experiment 2)

Condition	Word Conversion	Other Nonword	Letter Stripping	Other Letters
Repeated	.39	.22	.18	.21
Unrepeated	.54	.29	0	.17

The results of Experiment 2 are inconsistent with similarity inhibition models of RB: Even though the 3-letter stimuli shared a greater proportion of letters than the 5-letter stimuli, they did not demonstrate a greater degree of RB. The error analysis is also more consistent with a sublexical model; RB appears to selectively affect only the shared letters between W1 and W2.

General Discussion

Our previous investigations of RB between words of varying orthographic similarity (Harris & Morris, 1996, 1997; Morris & Harris, 1997) have suggested that RB between non-identical words is not a matter of similarity inhibition, but rather an emergent property of interactions at sublexical levels. In sublexical models of RB, only the shared letters between words are affected; the nonshared letters are unaffected and available for activating words. The two experiments presented in this paper, using word-nonword pairs as stimuli, provide additional evidence for the sublexical nature of RB. In Experiment 1, word-nonword pairs sharing all but one letter (Neighbor condition) showed greater RB than word-nonword pairs sharing only their first letters (Letter condition). This result is consistent with sublexical models if it is assumed that each shared letter, or letter cluster, can be differentially affected by the mechanism responsible for RB; it is also consistent with predictions of similarity inhibition models. However, in Experiment 2, three-letter word-nonword pairs sharing the same number of their letters as five-letter word-nonword pairs failed to show a greater amount of RB, in spite of their greater degree of orthographic similarity. This result is inconsistent with the predictions of similarity inhibition models. Furthermore, analysis of subjects' report errors in both experiments showed that shared letters are statistically less likely to appear in misreports of repeated conditions, compared to the same letters in unrepeated conditions. Non-shared letters, on the other hand, are *more* likely to appear in misreports of repeated conditions than are those same letters in unrepeated conditions.

Our finding of a sublexical locus for orthographic RB is also of interest in that it contradicts Kanwisher and Potter's assertion that RB occurs only at the level of representation being attended (Kanwisher & Potter, 1990). Although subjects are attending to and reporting words, RB is exhibited for letters within those words.

The results presented here are consistent with either single letters or letter clusters as the locus of sublexical RB effects; however, elsewhere, we have presented evidence in favor of a letter clusters view (Harris & Morris, 1997). We obtained greater amounts of RB between words sharing three consecutive letters (*await aware; genius gender*) than words sharing alternating letters (*above alone; gunmen gender*). Further exploration of the units involved in RB will shed light on general processes in word recognition and orthographic processing. These in turn will inform larger questions about subsymbolic processing and the perceptual-cognitive interface.

References

- Bavelier, D. (1994). Repetition blindness between visually different items: The case of pictures and words. *Cognition*, 51, 199-236.
- Bavelier, D., & Potter, M. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 134-147.
- Bavelier, D., & Jordan, M. I. (1992). A dynamical model of priming and repetition blindness. In C. L. Giles, S. J. Hanson, & J. D. Cowan (Eds.), *Advances in neural information processing systems*. San Mateo, CA: Morgan Kaufmann.
- 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, 1437-1455.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments & Computers*, 25, 257-271.
- Harris, C. L., & Morris, A. L. (1996). Letter clusters unite: Illusory words created by repetition blindness. *Poster presented at the Thirty-seventh Annual Meeting of the Psychonomic Society*, Chicago, Illinois.
- Harris, C. L., & Morris, A. L. (1997). The letter clusters theory of orthographic repetition blindness. *Poster presented at the Thirty-eighth Annual Meeting of the Psychonomic Society*, Philadelphia, Pennsylvania.
- Kanwisher, N. (1986). *Repetition blindness: Type recognition without token individuation*. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge.
- Kanwisher, N. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27, 117-143.
- Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 404-421.
- Kanwisher, N., & Potter, M. (1990). Repetition blindness: Levels of processing. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 30-47.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375-407.
- Morris, A. L., & Harris, C. L. (1997). A sublexical locus for repetition blindness: Evidence from illusory words. *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 519-524). Mahwah, NJ: Lawrence Erlbaum Associates.
- Mozer, M. C. (1983). Letter migration in word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 531-546.
- Mozer, M. C. (1991). *The perception of multiple objects: a connectionist approach*. Cambridge, MA: MIT Press.
- Park, J., and Kanwisher, N. (1994). Determinants of repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 500-519.
- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. E. (1996). Computational principles in quasi-regular domains: Understanding normal and impaired word reading. *Psychological Review*, 103, 56-115.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523-568.