Cognitive Psychology xxx (2008) xxx–xxx

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Repetition blindness: An emergent property of inter-item competition ☆,☆☆

Alison L. Morris [a,*], Mary L. Still [a], Catherine L. Caldwell-Harris [b]

[a] Department of Psychology, Iowa State University of Science and Technology, W112 Lagomarcino, Ames, IA 50011, USA
[b] Department of Psychology, Boston University, 64 Cummington Street, Boston, MA 02215, USA

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Repetition blindness
Repetition priming
Visual masking
Competition
Computational model
Simulation

ABSTRACT

Repeating an item in a brief or rapid display usually produces faster or more accurate identification of the item (repetition priming), but sometimes produces the opposite effect (repetition blindness). We present a theory of short-term repetition effects, the *competition hypothesis*, which explains these paradoxical outcomes. The central tenet of the theory is that repetition produces a representation with a higher signal-to-noise ratio but also produces a *disadvantage* in the representation's ability to compete with other items for access to awareness. A computational implementation of the competition hypothesis was developed to simulate standard findings in the RB literature and to generate novel predictions which were then tested in three experiments. Results from these experiments suggest that repetition effects emerge from competitive interactions between items and that these influences extend to adjacent, nonrepeated items in the display. The results also present challenges to existing theories of short-term repetition effects.

© 2008 Elsevier Inc. All rights reserved.

1. Introduction

Human observers are remarkably adept at identifying and reporting short lists of letters, words, or pictures displayed in rapid-serial-visual-presentation (RSVP; Forster, 1970). However, if two of the

☆ Experiments 1 and 2 of this article were presented at the 47th Annual Meeting of the Psychonomic Society, Houston, Texas, November, 2006.

☆☆ We wish to thank Michael Masson, Brad Wyble, and an anonymous reviewer for helpful comments on an earlier draft of this article. We also thank Miranda Atkinson, Nichole Stubbe, Megan Kienzle, Lisa McFadden, Haley Taylor, and Brandon Cornelius for their assistance in testing participants.

* Corresponding author. Fax: +1 515 294 6424.
E-mail address: almorris@iastate.edu (A.L. Morris).

0010-0285/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.cogpsych.2008.08.001

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

display items are identical, one of them often fails to be reported, and when observers are questioned further, they deny having seen the repetition. This robust phenomenon is known as repetition blindness (RB; Kanwisher, 1987). A similar effect occurs when identical items are shown in brief simultaneous displays (Bjork & Murray, 1977; Egeth & Santee, 1981; Luo & Caramazza, 1996; Mozer, 1989). RB does not depend on physical identity of the stimuli, as it has been found for words and letters of different case (Bavelier & Potter, 1992; Kanwisher, 1987; Marohn & Hochhaus, 1988) and pictures shown at different orientations or viewpoints (Kanwisher, Yin, & Wojciulik, 1999). RB has also been demonstrated for orthographically similar words such as *brush* and *crush* or *brace* and *crack*; in this case, the repeated items affected by RB appear to be letter sequences rather than whole words (Harris & Morris, 2000, 2001; Morris & Harris, 1999).

RB presents an interesting paradox in that repetition has been more often associated with facilitation in processing (the well-known repetition priming effect) than with impaired processing. Existing theories of RB have attempted to address this paradox, but none is able to account for all of the available data. We argue that existing theoretical accounts fall short of being able to explain the wide range of RB effects because the role of competition between items in rapid or brief displays has not been fully considered. Therefore, the purpose of this article is to present a new model—the *competition model*—that explains how competitive interactions result in short-term repetition effects. We first describe critical findings that pose problems for existing theories of RB and explain how the competition model was designed to account for these problematic findings. Next, we demonstrate that the competition model is able to simulate standard findings from RB experiments. After that, we test novel predictions made by the competition model. Finally, we discuss potential limitations of the model as well as possible modifications and extensions.

1.1. Current theoretical accounts of RB

Existing theories of RB can be classified into two main theoretical frameworks. The "activation" framework assumes that the ability to report an item's occurrence depends on activation of an abstract representation of the item. Theoretical accounts within this framework (e.g., Chun, 1997; Kanwisher, 1987, 1991; Kanwisher & Potter, 1989, 1990; Luo & Caramazza, 1995, 1996) emphasize the continuity of RB phenomena across stimuli and display formats. The "construction/attribution" framework, in contrast, eschews the notion of individual item representations that are activated or inhibited. Instead, perceptions of entire processing episodes are constructed on the basis of interactions between the stimuli, task, and context. Theoretical accounts of RB within this framework (e.g., Masson, 2004; Masson, Caldwell, & Whittlesea, 2000; Whittlesea & Masson, 2005) emphasize the variability of RB when different aspects of the processing episode are changed.

The major activation accounts of RB are the *type refractoriness hypothesis* (Luo & Caramazza, 1995, 1996), and the *token individuation hypothesis* (Chun, 1997; Kanwisher, 1987, 1991; Kanwisher & Potter, 1989, 1990). Both the type refractoriness and token individuation accounts assume the existence of *type* nodes, which are abstract (e.g., case-independent) representations. According to the type refractoriness hypothesis, once an item's type node is activated above its recognition threshold, further activation is briefly suppressed, analogous to a refractory period for a neuron. If a second item of the same type is displayed before the type node returns to baseline, the probability of the repeated item's being recognized is reduced. This assumption predicts that the major factor in determining the magnitude of RB should be the amount of time between encoding of the repeated items (Luo & Caramazza, 1996).

Whereas the type refractoriness hypothesis assumes that the type node is activated only once in the presence of two identical stimuli, the token individuation hypothesis assumes that the type node is re-activated by a repeated stimulus. However, the repetition is not consciously perceived because the re-activated type is not bound to a second spatiotemporal "token", or event representation (Chun, 1997; Kanwisher, 1987, 1991; Kanwisher & Potter, 1989, 1990). The mechanism responsible for the failure to "individuate" a second token for a repeated type has yet to be clearly specified. Candidate mechanisms have included a refractory period for token individuation (Kanwisher & Potter, 1989) or a token-minimization heuristic designed to avoid erroneously assigning two tokens to one event, as when one makes a saccade (Kanwisher, 1987).

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

Data challenging the idea that RB occurs as a result of temporal limitations in either type activation or token individuation can be found in an experiment by Whittlesea and Masson (2005, Experiment 1). In their experiment, participants viewed RSVP streams containing two critical words; for half of the trials, these words were identical. What also varied from trial to trial was whether or not there were other (filler) items in the display, and what these items, if present, consisted of. In the "blank" condition, the two critical words were shown for 120 ms each with a 120-ms blank interval in between. In three other conditions, the critical words occupied serial positions 2 and 4 in a five-item RSVP stream, and positions 1, 3, and 5 were occupied by different types of filler items: symbol strings, the word *WHITE*, or three different words. The task was simply to indicate whether or not a word (other than *WHITE*) was repeated within the stream.[1] Repetition detection was nearly 100% in the "blank" condition, but accuracy decreased for the other three conditions; participants correctly detected 78% of repetitions in the symbol-string condition, 58% in the *WHITE* condition, and only 10% in the different-words condition (false alarms were low in all conditions).

From these findings, Whittlesea and Masson (2005) concluded that RB cannot simply be the result of a temporal limitation on token individuation or type activation. They advocated abandoning the activation approach in favor of the construction/attribution framework, which argues that conscious experiences are the result of production and evaluation processes. For example, when a list of words is presented in RSVP, a mental event (a word coming to mind) may be produced. However, report of the word as having occurred in the list also requires attribution of the mental event to a particular source. Under this type of account, RB can occur when the mental event associated with the repeated word is misattributed to some other source, such as an identical or similar word that occurred earlier in the list (Masson, 2004).

Whereas the type refractoriness and token individuation hypotheses face difficulties in accounting for Whittlesea and Masson's (2005) findings, the results of other experiments are not easily explained by the construction/attribution hypothesis. Luo and Caramazza (1996, Experiment 1) examined the probability of report of letters in spatially distributed RSVP and simultaneous displays (the latter of which they called BSVP, for "brief simultaneous visual presentation"). In each of their displays, five letters and three symbols appeared in eight designated positions around an imaginary circle with the five letters appearing in adjacent positions. Participants were instructed to report items in clockwise order. On half of the trials, two of the letters were identical (in either the second and fourth or the third and fifth positions, clockwise). Critical letters were designated C1 and C2. On repeated trials, C1 and C2 were jointly reported at a lower rate than on nonrepeated trials (RB). Of particular interest was the finding that the letters adjacent to C2 were reported with greater accuracy in the repeated condition than in the nonrepeated condition; i.e., when the 2nd and 4th letters were identical, report of the 3rd and 5th letters increased, and when the 3rd and 5th letters were identical, report of the 4th letter increased (although the latter effect was not significant in the BSVP condition). Luo and Caramazza suggested that the finding of increased report of letters adjacent to C2 in the repeated condition is consistent with the type refractoriness hypothesis: When a repeated C2 is not recognized, it does not interfere with adjacent items entering a memory buffer. While construction/attribution accounts provide a reasonable explanation of why a repeated item might fail to be reported from an RSVP list, it is not clear how they would account for the finding of increased report of letters adjacent to the second of the repeated letters.

The failure of existing theoretical explanations of RB to adequately account for important empirical findings stems from an under-appreciation of the role of inter-item competition in producing the RB effect. The results of both the Whittlesea and Masson (2005) and Luo and Caramazza (1996) experiments can be readily explained if one adopts two assumptions. The first assumption is that items shown in close temporal proximity, whether targets, pattern masks, or filler items, must compete for access to conscious awareness (Enns & Di Lollo, 2000). The second assumption is that repeated items compete less effectively than do nonrepeated items. On this view, RB occurs because a repeated item is often out-competed by other items in the display; thus, viewers are more often "blind" to a

[1] Although RB is most often studied using a "full-report" paradigm (in which all of the stimuli in an RSVP stream or brief simultaneous display are to be reported), it can also be demonstrated in a repetition–detection task (in which the participant simply reports whether or not a repetition occurred).

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

repeated item than to a nonrepeated item. Adopting these assumptions leads to the prediction that the size of the RB effect should be magnified when competition between items is increased. Thus, greater RB is predicted in the Whittlesea and Masson experiment when the filler items are words than when they are symbol strings; although both words and symbol strings can induce competition for access to awareness, word fillers compete more effectively against the repeated target items, and therefore tend to increase RB. On the other hand, using the repeated word *WHITE* as the filler word is expected to *decrease* RB relative to the different-words condition, because repeated filler items compete against the targets less effectively than do nonrepeated filler items.

Competition between display items is also apparent in the results of Luo and Caramazza's (1996) experiment, in that decreased report of the repeated letter was associated with increased report of letters adjacent to the repeated letter. This finding is also consistent with the idea that repeated items are less effective competitors than nonrepeated items for access to awareness. But what exactly is it about repeated items that makes them less effective competitors? To answer this question, we turn to some recent findings in neurophysiology.

1.2. Neural correlates of stimulus repetition and conscious awareness

Recent investigations of the neural correlates of stimulus repetition have generally found that the total amount of activation associated with a repeated stimulus is smaller than that for a nonrepeated stimulus. In fMRI studies with human subjects, stimulus repetition typically results in a smaller area of activation (e.g., Buckner et al., 1998; Dehaene et al., 2001). Electrophysiological recordings from inferotemporal cortex in macaque monkeys revealed that some of the cells participating in the initial response to a stimulus decrease their firing rates when the stimulus is repeated (Baylis & Rolls, 1987; Miller & Desimone, 1994; Miller, Li, & Desimone, 1993). Furthermore, the particular cells that show suppression are not critical for representing the stimulus, while the activity of a small number of cells increases with repetition (Ringo, 1996). The selective suppression of non-critical cells coupled with the selective enhancement of others has been referred to as "sharpening" the representation (Desimone, 1996; Wiggs & Martin, 1998). Sharpened representations have two critical characteristics: The signal-to-noise ratio of the representation is increased (i.e., the activation of critical cells is increased relative to that of non-critical cells) and the overall activation associated with the representation is smaller. We suggest that what some RB theorists call the *type* representation is analogous to the representations that are subject to sharpening upon repetition; e.g., in a perceptual identification task with a briefly presented stimulus, a sharpened representation would be expected to produce more accurate identification. What remains to be described is how sharpening could be associated with priming in certain paradigms, but RB in others such as RSVP. Why would a sharpened type representation result in a failure to consciously perceive the stimulus?

A possible answer can be found by examining the neural correlates of perceptual awareness. Several investigators have suggested that sustained activation for several hundred milliseconds may be necessary for a representation to reach awareness. Subramaniam, Biederman, and Madigan (2000), citing findings from Tovée and Rolls (1995), suggested that most of the perceptual information from objects and faces is encoded in inferior temporal cells during the first 50 ms of firing, but sustained firing for an additional 300 ms is required for the percept to become stable enough to guide action and memory (see also Tononi & Edelman, 1998; Ward, 2003). Another factor that has been linked to perceptual awareness is synchronous activation of distributed cortical areas. For example, Dehaene et al. (2001) used fMRI and ERP methods to examine differences between the brain's responses to masked and visible words presented in RSVP. Briefly presented words (which could not be identified by the participants) were associated with reduced activation compared to unmasked (visible) words; in addition, masked words did not elicit the same correlated and distributed pattern of activation evoked by unmasked words. Furthermore, a large P300 component, generally associated with synchronous activation of distributed cortical areas, was observed only for visible words. Dehaene and Naccache (2001) suggested that neural activity must be sustained for a sufficient duration to activate a closed loop involving a wider neural network. Weak neural activity associated with brief, masked presentations of a stimulus will often fail to reach this "consciousness threshold"; however, as stimulus duration in-

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

creases and neural activity accumulates, the activation can become self-sustaining, and the stimulus representation gains access to awareness.

We now return to the question of why a more efficient or sharpened neural representation sometimes results in a failure to consciously perceive the associated stimulus: If a repeated stimulus generates less summed neural activity than a nonrepeated stimulus, then the representation of a briefly presented repeated stimulus would be less likely to exceed the consciousness threshold and access a self-sustaining state. Such a system would demonstrate a preference for novelty, in that the noisier representations associated with nonrepeated or less frequently encountered stimuli would provide them with preferential access to awareness (see Desimone, Miller, Chelazzi, & Lueschow, 1995, for similar ideas). However, the proposal that repeated items are associated with reduced activation is not sufficient to explain RB. In order to fully account for the paradoxical nature of repetition effects, we must also consider the role of competition in access to conscious awareness.

1.3. The competition hypothesis

The competition hypothesis proposes that sharpened stimulus representations, usually associated with priming, can instead result in RB when multiple neural representations are competing for access to awareness. RB is most often studied using RSVP, wherein each stimulus is presented for a brief period of time (generally less than 150 ms) and then is replaced by a subsequent stimulus. Each stimulus is therefore affected by both forward and backward masking. These are proposed to have different effects.[2] Forward masking from a visual stimulus adds noise to the representations of subsequently displayed stimuli (Breitmeyer, 1984). For example, when the letters *A*, *B*, and *C* are presented in RSVP, *A*'s neural representation will merge to some extent with *B*'s, adding noise to *B*'s representation; in addition, some of *A*'s representation (along with *B*'s) adds noise to the representation of *C*. The successive buildup of neural activity within the recognition system results in degraded stimulus representations (Desimone, 1996). However, if an item is repeated, as in display of *ABA*, the representation of the second *A* will suffer less forward interference than would a nonrepeated item *C*, because the two instances of *A* have similar representations; that is, the representation of a repeated *A* has an increased signal-to-noise ratio relative to that of a nonrepeated *A*. In addition, repetition of *A* tends to result in a smaller amount of total activation associated with the second *A* (a sharpened representation), as explained in the next section and illustrated in Fig. 1.

Whereas forward masking results in activation from previous stimuli being added to the representations of subsequent stimuli, backward masking does not add activation from later stimuli to the representations of previous stimuli. However, backward masking can interfere with identification of a stimulus' type by limiting processing time. Forward and backward masking also induce competition between stimuli for access to conscious awareness (Enns & Di Lollo, 2000). Which stimuli compete depends on two factors: processing time required for access to awareness and stimulus display rate. It has been suggested that sustained activation for 300–400 ms is required for a representation to access awareness (Subramaniam et al., 2000). But in most RSVP paradigms, new stimuli appear every 100–130 ms; thus, in order to enter awareness, a type representation must undergo additional activation beyond the time it is displayed. Because the pattern of activation at the type nodes is overwritten by each successive stimulus, this additional activation must take place at a different level of the system, known as the *event buffer*.

The visual system is sensitive to the abrupt changes in luminance which delineate separate events in the RSVP stream (Yantis & Jonides, 1984). Temporally distinct events are implemented in the competition hypothesis as follows. For each new stimulus, the current pattern of activation at the type nodes is continuously updated in the event buffer until a new stimulus appears, at which time the current event representation begins to gradually decay and the new event representation (occupying a separate set of nodes) begins updating. Multiple representations can thus be maintained in the event buffer; however, only one representation can enter awareness from the buffer at a time, and the ability

[2] Although masking effects occur at a number of levels in visual processing, in this article we are most concerned with "higher-level" masking effects, i.e., masking at the level of type representations.

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

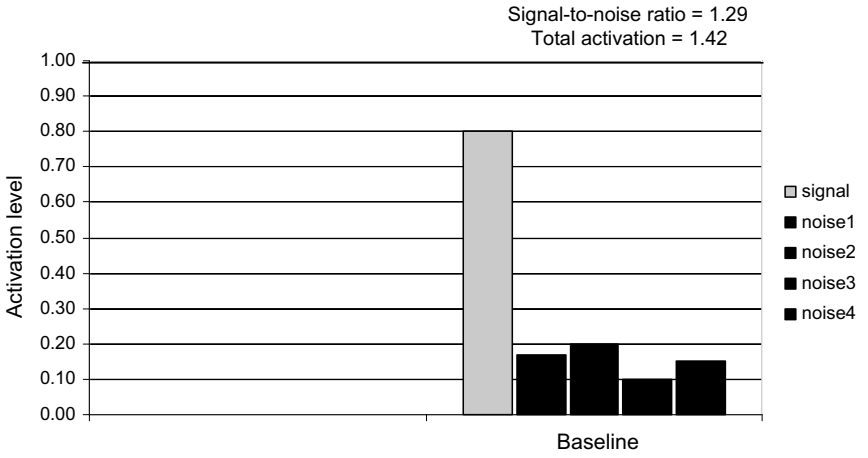

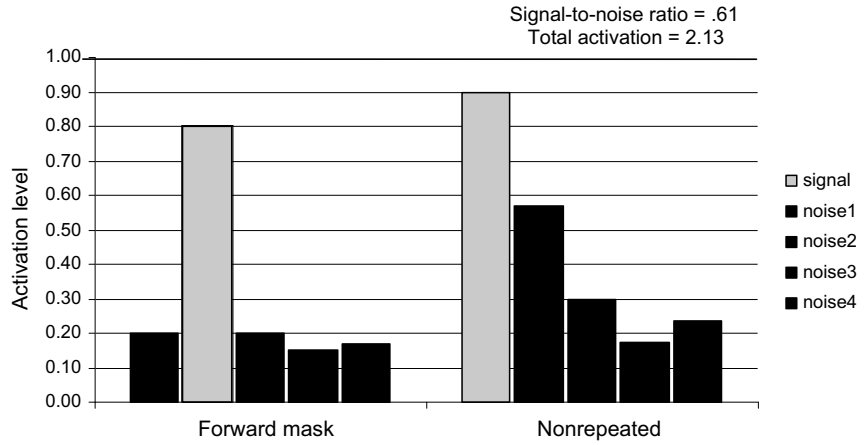

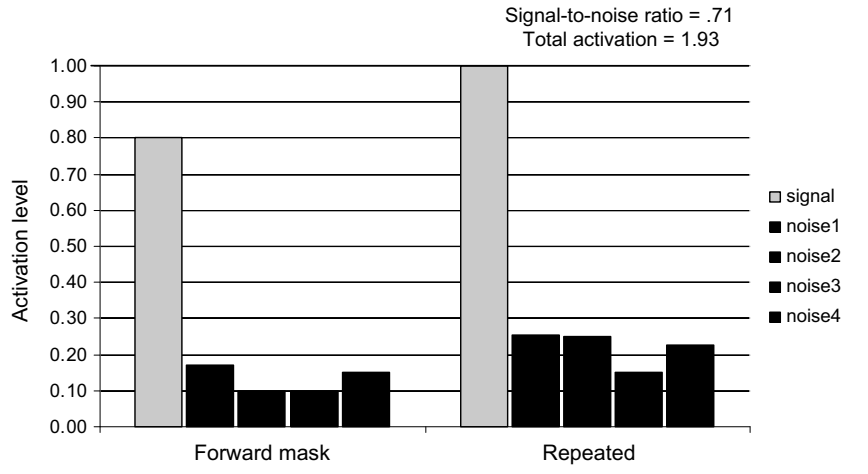

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

Fig. 1. The top panel shows encoding of the letter *A* without the influence of forward masking (baseline). Because of the brief display used in RSVP, the stimulus representation is part "signal" and part "noise". The middle panel shows the effects of forward masking from the letter *B* on the representation for the letter *A*. Decaying activation from the forward mask adds noise to each of the nodes in the representation for the letter *A*, decreasing the identifiability of the letter's type representation (represented as the signal-to-noise ratio) but increasing the total activation associated with the stimulus compared to baseline. The bottom panel shows the effects of forward masking from the letter *A*. Noise is added to each node, but because the activation patterns of the letter *A* and its forward mask *A* are similar, the identifiability of the type representation for the second letter *A* is less degraded (relative to baseline) than when *B* is the forward mask. However, because the signal node cannot exceed 1.0, the total activation associated with the repeated *A* is decreased relative to that for the nonrepeated *A*.

of a representation to stay active in the buffer is time-limited. Under typical RSVP conditions, the event buffer will contain two representations at a time, (e.g., item $N - 1$ and item N, or item N and item $N + 1$). This means that a stimulus' representation can only enter awareness if it wins the competition against the preceding or following stimulus. For example, suppose an RSVP series consisting of the letters *ABCDE* is displayed. *A*, *B*, and *C* can all access awareness if *A* wins the competition against *B*, *B* wins the competition against *C*, and *C* wins the competition against *D*. A stimulus fails to access awareness if it loses the competition against the representations of *both* the preceding and following items (for example, if *B* loses to both *A* and *C*, *B* fails to access awareness). Finally, if successive stimuli are displayed slowly enough, all stimulus representations can access awareness because each can undergo uninterrupted, sustained activation.[3]

What determines the winner of the competition between representations? According to the competition hypothesis, winning the competition for access to awareness depends on the amount of summed neural activity associated with each stimulus representation: Representations with larger activity typically out-compete representations with smaller activity. Repeated items, because of the smaller amount of neural activity associated with them, will generally be at a competitive disadvantage relative to nonrepeated items. Thus, although repeated items have an advantage over nonrepeated items in the form of sharper stimulus representations (which enables them to be more easily identified by the type recognition system), repeated items are *less* likely to access awareness than nonrepeated items.

Finally, because inter-item competition is based on summed neural activity, it can also be biased by attentional activation. In the absence of attentional bias, competition from backward masking is stronger than competition from forward masking. As an example, if *B* is displayed immediately following *A*, some of *A*'s activation is added to *B*'s representation, giving *B* a competitive advantage over *A*—with the result that *B* may be the only letter perceived (as occurs in masked priming). However, if more attentional activation is added to *A*, backward masking from *B* will be decreased (Enns & Di Lollo, 2000). In sum, access to awareness is an emergent property of inter-item competition based on summed activation, which in turn is influenced by repetition status and temporal attention.

1.4. A computational model

To examine the predictions of the competition hypothesis in detail, we implemented a simple model of the RSVP task using Matlab. Our strategy in developing the model was to employ the fewest assumptions necessary for simulating a variety of findings from RB experiments. The model simulates report of letters from a five-letter RSVP stream. According to the competition hypothesis, there are two requirements for a letter to be reported correctly: It must be identified correctly within the type recognition system, and it must gain access to awareness by out-competing other items in the display. The components of the implemented model include the output of a type recognition system and added activation from temporal attention.

In the competition model, types are distributed representations consisting of signal and noise. Each single letter is represented by a different pattern of activation across five nodes. For example, the letter *A* is represented by [10000], *B* by [01000], etc., such that for each letter, one node is a "signal" node

[3] The assumption that items compete in pairs in the competition model is a simplifying assumption based on the exposure durations typical of RSVP and the assumption that the buffer is time-limited. If much shorter exposure durations are used, more than two items at a time might be expected to compete for access to awareness.

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

and the others are "noise" nodes. Random activation is added to both signal and noise nodes to simulate the brief stimulus presentations typical of RSVP (complete specification of the implementation appears in Appendix A). This means that a five-letter RSVP stream is input to the model as a sequence of noisy letter representations. The distributed nature of the type representations in the competition model enables simulation of forward masking via slow decay at the type recognition-system output nodes, such that the type activation for each letter is partly a function of the letters preceding it. The firing rate at each type recognition-system output node (A_j) is given by

$$A_j(t) = A_i(t) + (1 - d)A_j(t - 1) \quad (1)$$

where t represents the serial position of the letter within the RSVP stream, A_i is the input activation to the node, and d represents a decay parameter. The firing rate at an output node has a maximum value of 1.0. All output nodes are reset to 0 at the end of each trial. The type recognition system outputs the correct letter when the activation at the signal node exceeds the activation at each of the other nodes by a criterion c.

When each new letter is input, the current pattern of activation across the output nodes (from the prior letter) is copied to an event buffer, in which the event representation of the prior letter can undergo the sustained activation required to access awareness. The events within the buffer compete on the basis of the summed activation (signal plus noise nodes) for each event, and the representation with the greater amount of neural activity accesses awareness. Thus, if the summed activation of the first letter in an RSVP stream exceeds that of the second letter, the first letter accesses awareness, and the second letter must continue competing with the third letter. If the third letter also out-competes the second letter, the second letter will fail to access awareness; this occurs because of the time-limited nature of the event buffer. The final letter in the RSVP stream, which is followed by a symbol-string mask, can access awareness if its activation exceeds either that of the preceding letter or a fixed parameter m (which represents activation of the symbol mask). It is possible for a letter with an incorrect or ambiguous type representation to out-compete adjacent items and access awareness; in this case the letter will not be reported correctly.

A final component of the model is temporal attention, which biases the competition between RSVP stream items. Temporal attention is simulated by adding attentional activation to the sum of the activation across all output nodes for a given letter in the event buffer, so that total activation associated with a letter is given by

$$A_{\text{tot}}(t) = \sum_j A_j(t) + att(1 - ad)^{(t-1)} \quad (2)$$

where att represents the attention parameter and ad is an attentional decay parameter. Because competition for access to awareness depends on the total activation associated with a letter's representation (including attentional activation), letters receiving more temporal attention will have a competitive advantage. For example, if a large amount of attentional activation is applied at the beginning of the RSVP stream, and that activation gradually decays, the summed activation for the first letter will exceed that of the second letter, and the second letter will exceed that of the third, with competition between the letters increasing as the stream progresses. Optimal values for the att and ad parameters enable each letter in a short RSVP stream to access awareness by winning the competition against the letter immediately following it; large but non-optimal values of these parameters result in increased report of items at the beginning of the stream relative to those nearer the end.

Although in most RB experiments participants are instructed to report all the items displayed, in some experiments participants are required to report only the final item in the RSVP stream. According to the competition model, temporal attention and the rate of its decay may be varied strategically in order to favor items at the end as well as the beginning of the RSVP stream. Because temporal attention applied at the beginning of each trial gradually decays, items at the beginning of short RSVP streams generally have the highest total activation, with activation decreasing as serial position increases. This activation function favors items at the beginning and middle of the stream at the expense of those at the end. However, if the attentional decay is more rapid, the resulting activation function will be U-shaped; items at the beginning of the stream still out-compete successive items, but as activation from temporal attention decays, items in the middle of the stream begin to lose the ability to

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

overcome backward masking from the items that follow them. In this situation, the final item in the stream has a competitive advantage. Thus, when instructed to report all items in the order displayed, viewers will apply a slow rate of attentional decay, but when instructed to report the last item in the stream, attentional decay will be set at a higher level.

How does the model infer the temporal order of rapidly displayed items? An ancillary assumption of the competition model known as the *temporal order assumption* states that viewers may use the relative amounts of summed activation associated with successive stimuli to infer the order in which they appeared.[4] When viewers are presented with a short RSVP list that is to be reported in order, items appearing early will have greater activation than items appearing later; thus, when determining the order of two successively presented items, viewers will assume that the item with greater activation was presented first. However, if viewers are instructed to report the final item in a longer list, they will exploit the fact that applying less attention will increase the competitiveness of items at the end of the list. Under these circumstances, when making judgments about items near the end of the list, viewers will assume that an item associated with higher activation occurs *later* than an item associated with lower activation.

The competition model superficially resembles the token individuation hypothesis in that both theories utilize type representations and event representations; however, it is important to note that RB in the competition model does not represent a failure to bind repeated type representations to separate events. According to the competition model, event representations are formed for both repeated and nonrepeated items, but these representations must then compete with temporally adjacent items for access to awareness. Because of this competition, both repeated and nonrepeated items sometimes fail to access awareness. However, repeated items, by virtue of their smaller summed activation, fail to access awareness more often than do nonrepeated items. As described in the next section, this explanation for the RB effect is able to account for a wide range of existing findings in the literature.

2. Competition model simulations

As an initial test of the competition model, we simulated the results from Luo and Caramazza (1996) shown in Fig. 2a and b. As described previously, Luo and Caramazza displayed five letters (which we will designate L1, L2, L3, L4, and L5) in spatially distributed RSVP and BSVP, and found that when the critical letters were L2 and L4, RB for L4 was accompanied by increased report of L3 and L5. There were also two rather puzzling findings. The first was that the magnitude of RB was larger in spatial RSVP than in BSVP; the second was the finding of "backward RB", that is, RB in both paradigms was found for C1 as well as C2. Simulating Luo and Caramazza's results enables us to explore the dynamics of inter-item competition and may help to clarify the similarities and differences among standard RSVP, spatial RSVP, and BSVP paradigms.

One obvious difference among the different RB paradigms is that in standard RSVP all stimuli are displayed at the same spatial location, whereas in spatial RSVP and BSVP stimuli are displayed in different spatial locations. In addition, both standard RSVP and spatial RSVP items are displayed successively, whereas in BSVP all items are displayed simultaneously. As can be seen in Fig. 2a and b, the results from Luo and Caramazza's spatial RSVP and BSVP displays were strikingly similar, prompting them to argue that the difference between successive and simultaneous presentations "may be more apparent than real" (Luo & Caramazza, 1996, p. 100). In other words, even in simultaneous displays, activation of letter type nodes occurs sequentially.

The competition model provides a framework for understanding the similarities and differences between spatial RSVP and BSVP. BSVP resembles spatial RSVP because participants covertly shift spatial attention from item to item, imposing a sequence on type activation.[5] With sequential activation of type representations, forward and backward masking will occur at the type level even when the stimuli are displayed at different locations. Consequently, the competition hypothesis predicts that the same

[4] See Page and Norris (1998) and Reeves and Sperling (1986) and for similar ideas.

[5] With a display duration of only 300 ms, Luo and Caramazza's BSVP displays were too brief to permit eye movements to each item's location.

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

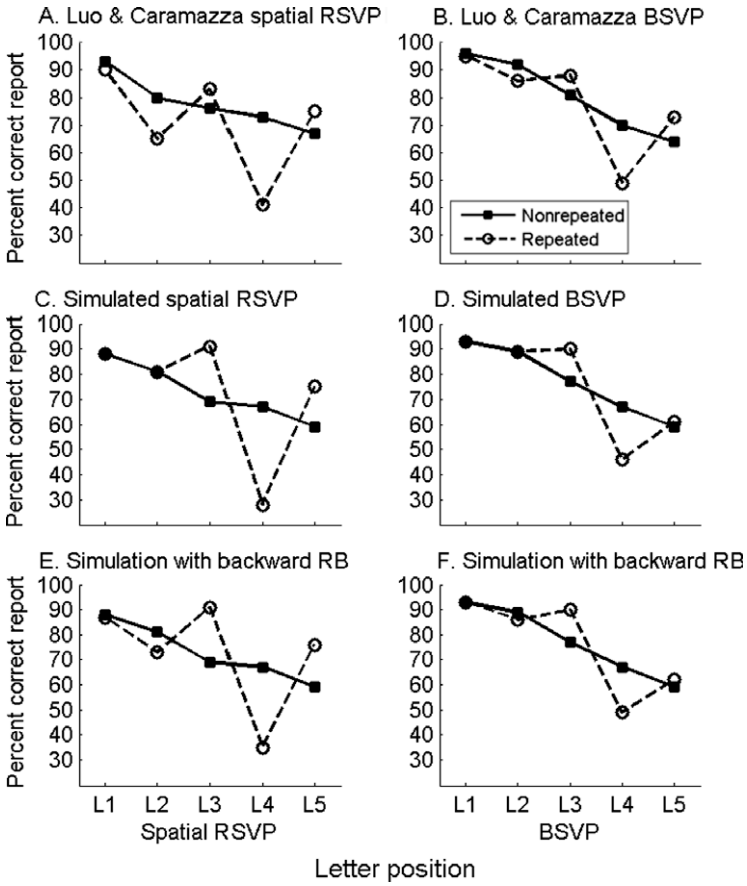

Fig. 2. (A and B) The results from Luo and Caramazza (1996). Competition model simulations of these results are shown in (C)–(F). Simulations shown in (E) and (F) include additional code used to calculate backward RB (see Appendix A for details).

types of competitive interactions found in RSVP should also be found in spatial RSVP and BSVP. But there should also be some small differences between spatial RSVP and BSVP, including a reduced RB effect in BSVP. In spatial RSVP, the sudden onset of each item will prompt a shift of attention to the new item. In BSVP, because all items onset at once, the rate that the individual items are sampled can vary from participant to participant; some participants may shift attention rapidly from item to item (making it similar to spatial RSVP) whereas others may shift attention more slowly. Slower shifting of attention means that activation at the recognition output nodes and attentional activation will both decay a bit more between stimulus presentations. In simulating spatial RSVP and BSVP, therefore, the values of the *d* (decay) and *ad* (attentional decay) parameters were increased slightly from that expected for standard RSVP and the decay rate was made more variable for BSVP (see Appendix A for additional details and parameter values). The slower progression from item to item in BSVP produces less forward masking at the type level, and less forward masking in turn produces less sharpening of the repeated C2 representation and therefore less RB. Results of the simulations are shown in Fig. 2c and d. As can be seen in the figures, RB for BSVP was smaller than that for spatial RSVP, just as it was in Luo and Caramazza's experiment. Thus, the major difference between BSVP and spatial RSVP was successfully simulated by the competition model.

However, our initial simulations failed to produce the "backward" RB found by Luo and Caramazza (1996). According to both the competition model and the type refractoriness hypothesis, RB is strictly a forward effect. Luo and Caramazza suggested that backward RB could occur on occasions when C2 is

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

Table 1
Simulated percent correct report of critical letters as a function of C1–C2 lag

	Lag 1		Lag 2	
Condition	C1	C2	C1	C2
Nonrepeated	100	67	100	61
Repeated	100	39	100	51
Nonrepeated–repeated		28		10

Note. C1 = first critical word; C2 = second critical word.

identified prior to C1. Although this is a reasonable explanation for BSVP paradigms, it is less amenable to RSVP displays. In RSVP it is unclear how type activation associated with C1 could occur later than that associated with C2, particularly when C1 and C2 are separated by another stimulus. Luo and Caramazza also considered the possibility that backward RB could be produced by migration errors; i.e., apparent "backward" RB can be produced if C1 is reported, but in the position of C2. For example, suppose that *ZQSQK* is displayed, and the participant successfully identifies the first *Q, S*, and *K* but reverses the order of *Q* and *S*, thus reporting *SQK*. In this case, C2 would be credited with the report of *Q*, and the trial would appear to be a possible instance of backward RB.

Adding the temporal order assumption to the competition model enables it to simulate backward RB. When participants are instructed to report items from short RSVP or BSVP displays in order, temporal attention is applied at the beginning of processing, and attentional decay is set so that total activation per item decreases with increasing serial position. Accordingly, the item with the higher level of overall activation is reported as having occurred first when two adjacent items both access awareness. Backward RB should occur when both L2 and L3 access awareness (L2, by out-competing L1, and L3, by out-competing L4), and L3's activation is higher than L2's; in that case, L2 would be reported as occurring after L3. Because of noise in the system, this situation has the potential to occur on a number of trials. Adding this assumption to the simulation produced a small amount of backward RB for both the spatial RSVP and BSVP conditions, as shown in Figs. 2e and f. Thus, the competition model demonstrates how backward RB can appear in the data when RB itself is strictly a forward effect.

We next present a series of simulations that illustrate how the competition model accounts for core findings in the RB literature, including: (1) the lag effect; (2) the relationship between encoding of critical items and RB; and (3) the relationship between repetition priming and RB (additional details and parameter values for each simulation can be found in Appendix A).

2.1. The lag effect

One of the most robust findings in the RB literature is the lag effect, whereby the magnitude of RB decreases as the number of items intervening between the critical items increases (e.g., Chun, 1997; Kanwisher, 1987; Park & Kanwisher, 1994). According to the competition model, increasing the lag between C1 and C2 decreases the influence of forward masking from C1 on C2; this decreases the amount of sharpening of C2's type representation, leading to an increase in its competitiveness and a decrease in RB. We simulated the lag effect using a five-letter standard RSVP stream; for lag 1, C1 was L1 and C2 was L3; for lag 2, C1 was L1 and C2 was L4. C1 and C2 were either identical (repeated condition) or non-identical (nonrepeated condition). Results of the simulation are shown in Table 1; RB is indexed by the difference in report of C2 in the repeated and nonrepeated conditions. As can be seen in the table, RB decreases as lag increases, consistent with the standard findings.

2.2. Encoding effectiveness and RB

Luo and Caramazza (1995) varied the exposure duration of C1 from 25 to 200 ms and found that increasing the encoding effectiveness of C1 by displaying it for a longer duration increased the mag-

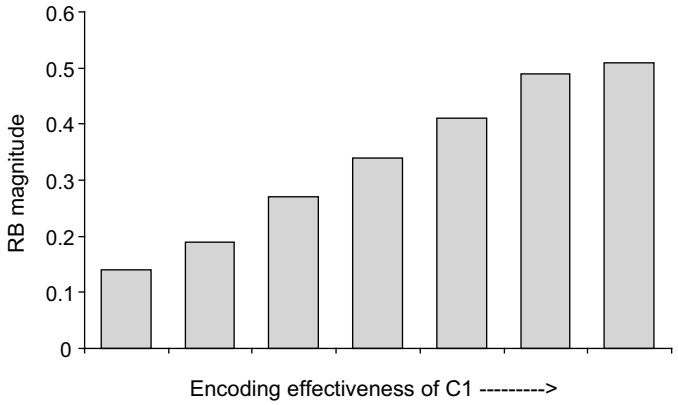

Fig. 3. Competition model simulation of the influence of encoding effectiveness of C1 on RB for C2. Encoding effectiveness was increased by increasing the signal-to-noise ratio of the type representation for C1. RB magnitude was calculated as the difference between correct report of C2 in the repeated and nonrepeated conditions.

nitude of RB for a repeated C2.[6] The competition model also predicts that a well-encoded C1 produces more RB for C2, because it produces greater sharpening in C2's representation. Encoding effectiveness can be improved in several ways, including increased stimulus duration, increased stimulus contrast, and reduced forward masking. In the model, a well-encoded C1 is represented by high activation at the signal node and low activation at the noise nodes (high signal-to-noise ratio). A poorly encoded C1 is represented by lower activation at the signal node and higher activation at the noise nodes. When C1 is poorly encoded, forward masking will increase the signal node activation for an identical C2, but because the signal node activation for C1 is not sufficiently high, the signal node activation for C2 will not reach ceiling. Therefore, although the signal-to-noise ratio of C2's representation is increased, its overall activation is not decreased relative to that of a nonrepeated C2. Consequently, there will be no competitive disadvantage for the repeated item and therefore no RB.

We simulated the influence of encoding effectiveness by gradually decreasing the amount of input noise; this has the effect of increasing the signal-to-noise ratio in the type representation. As shown in Fig. 3, increased encoding effectiveness of C1 was associated with an increase in RB for C2, consistent with the results from Luo and Caramazza (1995).

An interaction of encoding effectiveness manipulations with overall performance levels may also explain conflicting findings concerning the influence of word frequency on the magnitude of RB. Bavelier, Prasada, and Segui (1994) displayed RSVP sequences of three words followed by a symbol string as a final mask. In their first experiment, RB for orthographic neighbors (e.g., *cent* and *dent*) was greater when C1 was a high-frequency word. This result is consistent with the idea that, on average, high-frequency words are better-encoded than low-frequency words, and the increased encoding effectiveness results in greater RB for C2. However, in a second experiment using critical words that were identical (e.g., *down* and *down*), C1 frequency had no influence on the magnitude of RB. Although one difference between these experiments was the type of materials used (orthographic neighbors vs. identical words), a second difference was the amount of RB obtained: The average size of the RB effect at lag 1 was much larger using identical words in Experiment 2 (50%) than when using orthographic neighbors in Experiment 1 (21%). Although it may seem counterintuitive that the effects of a variable such as word frequency would be *less* evident when RB is larger, the competition model can help to illustrate why this should be the case.

We simulated the effects of word frequency when RB is large, medium, or small using three-item RSVP streams in which the first and last items were identical on half of the trials. Larger or smaller amounts of RB were produced by varying the amount of input noise. Lower-frequency words were

[6] See Section 7.1.4 for additional information concerning items with longer durations.

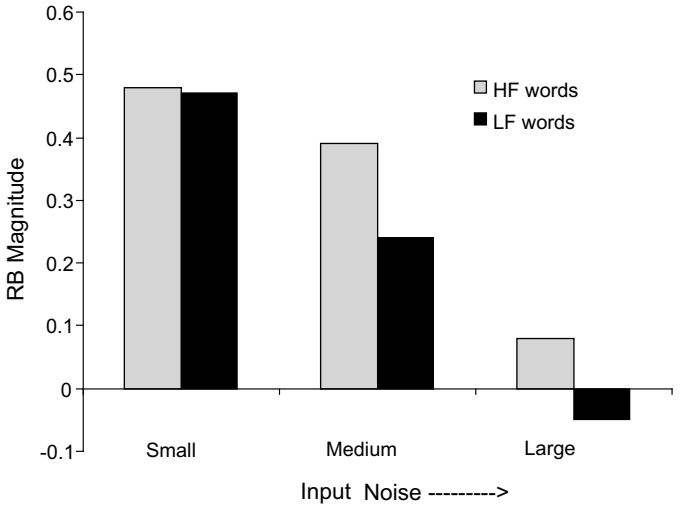

Fig. 4. Simulation of the effects of word frequency on RB magnitude at three different levels of input noise. When input noise is small, words are well-encoded and the magnitude of RB is large; as input noise is increased, the magnitude of RB decreases.

simulated by adding a small amount of additional noise to the type representations. The added noise decreased the signal-to-noise ratios for low-frequency words, making their types less identifiable, but increased their competitiveness (because adding noise increases total activation). Fig. 4 shows the predicted behavioral data; when RB is large, word frequency has little effect, but as the overall level of RB decreases, RB is greater for high-frequency words than for low-frequency words.

Fig. 5 demonstrates why the effects of word frequency are obscured when RB is particularly large. In the simulation, the activation associated with a repeated C2 was nearly always lower than that of the item preceding it; however, C2 could be reported if its total activation exceeded that of the symbol mask. As shown in Fig. 5a, when the total activation for a repeated C2 is much lower than that of the symbol mask (large RB), adding extra activation to the input has little effect; although a low-frequency C2 loses the competition with the mask by a smaller magnitude, it still loses. However, when the difference between C2's activation and that of the mask is smaller (Fig. 5b, medium RB, and Fig. 5c, small RB), a low-frequency C2 can more often out-compete the mask and be reported. Thus, when the magnitude of RB is smaller, frequency differences are more likely to be observed.

2.3. The relationship between short-term priming and RB

Examination of competition dynamics using the competition model also explains how repetition in rapid displays can lead to priming in some situations and RB in others. In standard RB paradigms, items are usually displayed for more than 100 ms, which is generally sufficient to ensure a high identification rate within the type recognition system; limitations on the number of items correctly reported stem largely from competition for access to awareness. However, a number of experiments have investigated the accuracy of naming a briefly displayed word (67 ms or less) at the end of an RSVP stream (e.g., Kanwisher, 1987, Experiment 3). According to the token individuation hypothesis, RB will only occur for a repeated C2 if it needs to be individuated from C1; therefore, if the task is to report only the final word in an RSVP stream (C2), C1 will not be individuated, and the result should be priming rather than RB for C2. The competition hypothesis also predicts this result, but not because of a change in processing strategy. When C2 is displayed for 100 ms or more, forward masking from an identical C1 drives its signal node to ceiling, limiting the growth of total activation. However, when the exposure duration for C2 is short, forward masking from an identical C1 will improve C2's signal-to-noise ratio without producing a ceiling effect. Thus, when C2 is poorly encoded, the total activation for

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001

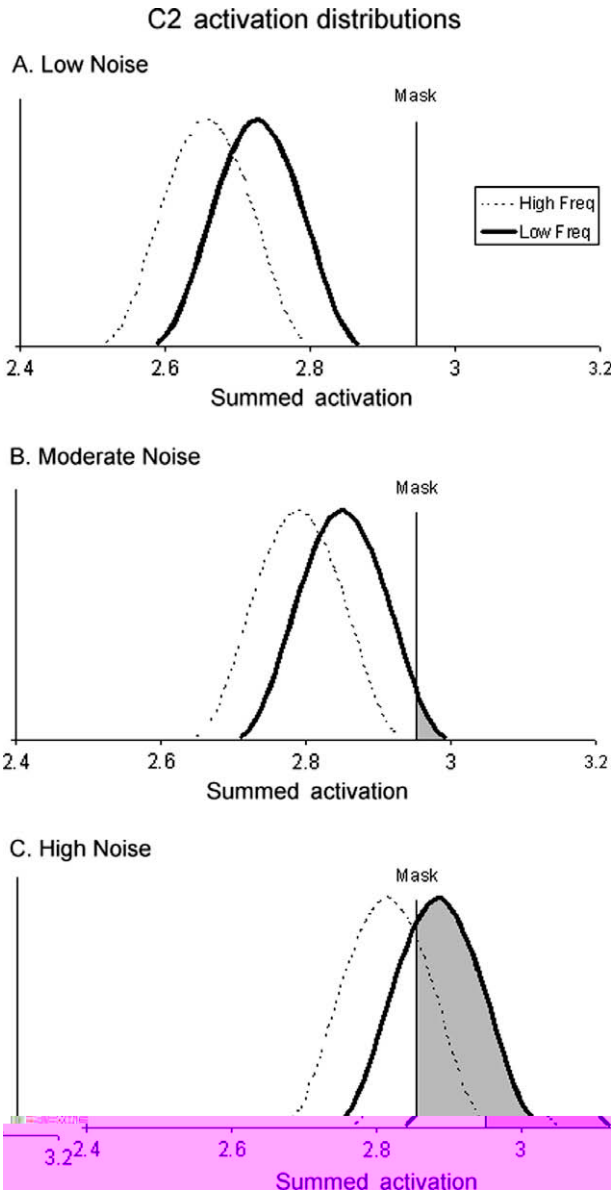

Fig. 5. Simulated high- and low-frequency C2 activation distributions relative to mask activation for three levels of input noise in the competition model. At low levels of input noise (top panel), the activation levels of both high- and low-frequency words are consistently exceeded by that of the mask, producing similarly large amounts of RB for both types of words. However, at moderate and high noise levels (B and C), the activation levels of low-frequency words exceed that of the mask more often than do the activation levels of high-frequency words, producing a decrease in RB for low-frequency words relative to that for high-frequency words.

a repeated C2 will not be decreased relative to a nonrepeated C2. This situation produces priming rather than RB. In Kanwisher's experiment, correct report of the final word averaged only 41% in the nonrepeated condition, improving to 59% in the repeated condition. We simulated these results

Please cite this article in press as: Morris, A. L., et al. Repetition blindness: An emergent property ... *Cognitive Psychology* (2008), doi:10.1016/j.cogpsych.2008.08.001