

High-level effects of masking on perceptual identification

Adam N. Sanborn^{a,*}, Kenneth J. Malmberg^b, Richard M. Shiffrin^a

^a Department of Psychology, Indiana University – Bloomington, 1101 E. 10th Street, Bloomington, IN 47405, USA

^b Iowa State University, IA, USA

Received 16 June 2003; received in revised form 6 January 2004

Abstract

The extent to which visual form versus higher-level information is used to identify briefly flashed words is assessed in a perceptual identification task. In this task, a word is briefly flashed, post-masked, and a decision is made between two alternatives. The availability of visual (e.g., case or color) and higher-level information (e.g., abstract letter codes, phonology, and meaning) was manipulated by varying the information that discriminates the alternatives. Performance was better with higher-level than with visual information when pattern post-masks were used, but the reverse occurred without masking. The authors conclude that both higher-level and visual-form information can be used to identify words with the strategy depending on the information available at the time of the choice.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Cognition; Form perception; Masking; Reading

1. Introduction

Cognition is often characterized as operating on a mixture of perceptually-based lower-level information and cognitively-based higher-level information (e.g., Broadbent, 1958, 1971). Many results show that two or more levels of information are available and used to identify or detect even relatively simple stimuli (e.g., Moray, 1959; Reicher, 1969; Sperling, 1960; Treisman, 1962). A classical question is the extent to which the performance of various tasks depends on the differential utilization of these levels of encoded information. The present concern is the extent to which high- versus low-level information is used to identify words presented visually near the threshold of detection, and how this weighting changes between masking conditions.

Investigations of this and similar issues often attempt to limit attentional resources available for processing low- versus high-level information (termed capacity limitations), or attempt to degrade low- versus high-level information (termed data or information limitations).

The present experiments follow the latter path by using a visually based post-masking procedure: A word is flashed briefly and followed first by a pattern mask and then by two choice words, one of which must be judged to be the flashed target. The two-alternative forced-choice procedure (i.e., 2AFC procedure) has an advantage over a simpler naming task because it allows the manipulation of the similarity of the choices, thereby allowing assessment of the utility of different stimulus aspects for identification. The 2AFC procedure also affords the possibility of deciding whether differences in performance are due to stimulus sensitivity or decision bias (cf. Broadbent, 1967; Green & Swets, 1966; Huber, Shiffrin, Lyle, & Ruys, 2001; McKoon & Ratcliff, 1995; Schooler, Shiffrin, & Raaijmakers, 2001; Zeelenberg, Wagenmakers, & Raaijmakers, 2002).

The perception literature is filled with a century's worth of demonstrations of the roles of both low- and high-level information. Examples in the domain of auditory word perception at threshold include the McGurk effect showing an interaction of auditory and visual information on speech perception (McGurk & MacDonald, 1976). Here, interference must be due to high-level codes because the low-level information is presented in different modalities. In the visual domain, letter detection improves when the letter is embedded in

* Corresponding author. Tel.: +1-812-855-0626.

E-mail address: asanborn@indiana.edu (A.N. Sanborn).

a letter string forming a word as opposed to in a letter string forming a non-word (e.g., the ‘word superiority’ effect, Reicher, 1969; Wheeler, 1970). The literature attempting to link the various empirical findings to models is large and complex, and we shall focus discussion upon the present paradigm after the results are presented.

2. Experiment 1

A simple model of word identification posits comparison of the visual form of the stimulus to the contents of the visual information in a mental lexicon (e.g., Morton, 1969). When a match is found between the form of the test word and the form in some lexical trace, the word is identified. Somewhat more complex models allow higher and more abstract levels of encoding at test, and comparison of that information to similar information stored in the lexical traces. These other levels include orthographic information, letter information, letter groups, phonological codes, and semantic codes (though semantic coding en route to word identification raises some additional issues).

All levels of information originate and stem from visual-form information, making it difficult to discriminate between models assuming identification based on different combinations of levels of coding. We decided to approach this difficult problem with the 2AFC technique because it allowed us to vary similarity of the choice words to each other. In what we shall term the ‘different-word’ condition, for instance, the word “BRAIN” might be flashed, masked, and followed by the choices “BRAIN” and “DRAIN”. In this situation, the alternative choices differ in both higher- (i.e., orthography, phonology, and meaning) and lower-level form information (i.e., the shape of the stimulus). In what we term the ‘different-case’ condition, the alternatives in this example are “BRAIN” and “brain”, which differ only in lower-level form information (due to the difference in case). If higher-level information is typically used to identify briefly flashed words, then performance should be worse in the different-case than in the different-word condition because the higher-level differences between the choices do not exist in the former. In the ‘both-different’ condition, the alternatives are “BRAIN” and “drain”. One might expect best performance in this condition because every level of encoding should distinguish the alternatives. Although these examples used ‘case’ as a distinguishing feature, we used a second condition with a color shift to increase the generality of the findings (Sperling, Wurst, & Lu, 1993). In the color conditions, upper and lower case were replaced by a color difference (all words were in upper-case).

2.1. Methods

2.1.1. Design and materials

Word identification was assessed using a 2AFC post-masking procedure. Each trial consisted, in temporal order, of: (1) a 400 ms central fixation of four dashes (----); (2) a 300 ms blank field; (3) central presentation of the target word (for a predetermined brief duration); (4) a 300 ms mask of seven @-signs (@@@@@@) that completely covered even the longest target word; (5) two alternative words (a target and a foil) presented to the left and right of fixation (see Fig. 1). The observer was asked to make an untimed choice of the alternative that matched the flashed target. Each experimental session consisted of 330 trials, with 1 s intertrial intervals. Trials were equally divided among five different flash durations, chosen on the basis of pre-testing to span the threshold range for each observer, giving 22 observations per flash duration per choice condition per session per observer. All observers completed 13 sessions. The data for psychometric functions were collected using the method of constant stimuli.

Pairs of words differing in a single letter were used as the target and foil choices. Three hundred and forty pairs were constructed from medium frequency words four to eight letters in length. Stimuli were presented on CRT monitors with refresh rates of 120 Hz. The visual angle of a presented word ranged from approximately 7° to 15°, with each word pair randomly assigned to a condition for each observer during each of 13 sessions. The target was randomly chosen from the pair selected

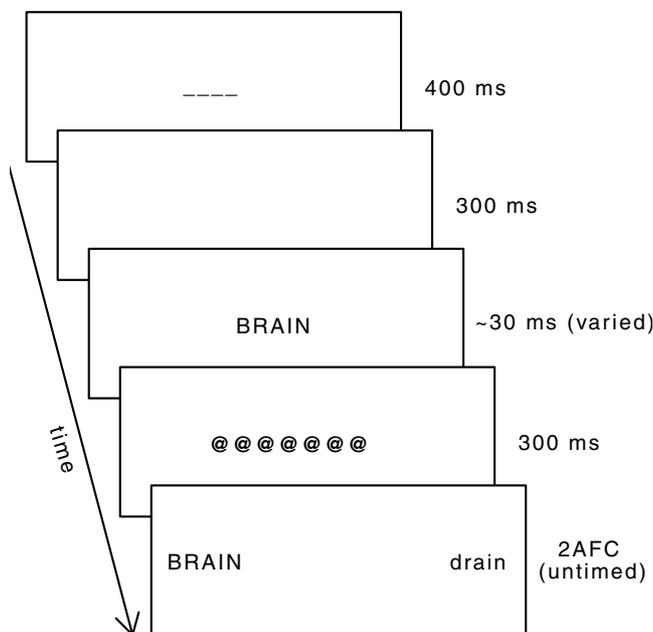


Fig. 1. Sequence of events in Experiment 1 for the case condition in which the two alternatives differed in both spelling and case (both different).

Table 1
Examples of stimuli and conditions from Experiment 1

Target	Left choice	Right choice	Condition
<i>(a) Case conditions</i>			
BRAIN	BRAIN	DRAIN	Different word
BRAIN	BRAIN	brain	Different case
BRAIN	BRAIN	drain	Both different
<i>(b) Color conditions</i>			
BRAIN [red]	BRAIN [red]	DRAIN [red]	Different word
BRAIN [red]	BRAIN [red]	BRAIN [green]	Different color
BRAIN [red]	BRAIN [red]	DRAIN [green]	Both different

A color name in brackets indicates the color in which the word appeared.

for a given trial, as was the left/right position of the target among the two alternatives. The case and color conditions described below were run in different sessions for a given observer.

Table 1 illustrates three of the 12 choice conditions for the case sessions (those for uppercase targets, and for target choice positioned left) and three of the 12 choice conditions for the color sessions (those for red targets and for target choice positioned left). Consider first the case conditions. In the different-word condition, the alternatives differed in spelling (and hence differed in orthography, phonology, meaning, etc.). In the different-case conditions, the alternatives differed in case only. In the both-different condition the choices differed in both case and spelling. In these case conditions the flashed word appeared in black on a white background in uppercase and lowercase equally often, and was followed by a mask of black @-signs on a white background. In the color conditions all words were in uppercase, and the case differences described above were replaced by color differences (red and green words), with the words presented on a gray background (RGB values of 255,0,0; 0,255,0; 127,127,127 respectively). The mask in the color conditions consisted of seven interleaved red and green @-signs (red always first). Within a session the order of the choice conditions was mixed, and randomly permuted so that each experimental session consisted of 110 trials of each choice condition.

2.1.2. Observers and procedure

Four Indiana University students were paid for their participation in 13 sessions each. There were five consecutive experimental sessions for case and five for color. Observers 1 and 3 participated in the case sessions first and then the color sessions; this order was reversed for observers 2 and 4. Sessions 1 and 2 were used to calibrate the five flash durations to span the accuracy range for the first set of sessions. Session 8 was used to recalibrate the flash durations for the other condition (case or color). The flash durations were allowed to differ among the three choice conditions tabled, but this

was unnecessary except for observer 4 in the color sessions, for whom the different-color condition required longer flash times. Observers were run on the same monitor every session to control for slight color imbalances across monitors. Word pairs were chosen randomly for the calibration and experimental trials, and no words were repeated across trials during any one session. Observers were instructed to focus attention upon the fixation position and the briefly flashed word that would follow, and to make their best guess concerning which choice word exactly matched the one flashed.

2.2. Results and discussion

Observers used higher-level cognitive information instead of lower-level case or color information (see Figs. 2 and 3). For each choice condition, cumulative Weibull distribution functions were fit to the mean accuracies as a function of flash duration using maximum likelihood (to avoid local minima, 10 random sets of start values for the parameters were used to fit each condition, and the best fit was chosen). The cumulative Weibull distribution was allowed to range from 0.5 (i.e., chance) to 0.99 (i.e., perfect performance minus a 0.01 probability of stimulus-independent error). Differences between conditions were assessed with likelihood ratio tests. The maximum likelihood goodness of fit for the three choice conditions fit with three functions (i.e., the full model) was compared to the maximum likelihood of more restricted models. The first restricted model (i.e., two curves) fit two Weibull functions to the data: one for the different-case (or different-color) condition and one to the other two conditions. The second restricted model (i.e., single curve) fit one function to the data from all three choice conditions. The likelihood ratio of nested models (i.e., $-2[\log(\text{restricted}) - \log(\text{full})]$) follows a χ^2 distribution with degrees of freedom equal to the difference in number of parameters. A significant likelihood ratio test (assumed to be $p < 0.05$) meant that the restricted model fit the data worse than the full model.

The data, fits, and likelihood ratio tests for each observer are shown in Figs. 2 and 3 for the case and color conditions respectively. The common result was that the two curve model did not fit worse than the full model, while the single curve model was inferior to the full model. This result held for all four observers for case and three of the four observers for color (excepting observer 1's color responses, for which the two curve model was worse than the full model). Thus, the two conditions that differed in spelling did not differ reliably from each other, but both were reliably superior to the condition that differed only in case or color.

The preference for high-level information is initially counterintuitive because it might have been expected that the easier discriminations would have been those

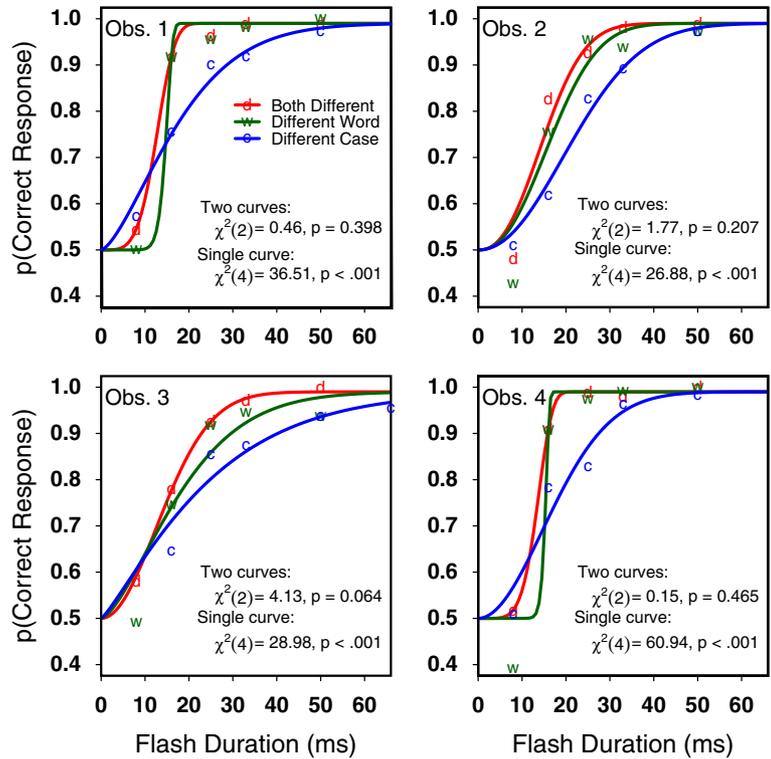


Fig. 2. Case condition's data shown by letters, Weibull fits shown by curves, and likelihood ratio tests (χ^2). Each panel is an observer. Conditions are shown in Table 1. Likelihood ratio tests are between the full model (curves shown) and restricted models (specified in the text). Numbers in parentheses are degrees of freedom for the likelihood ratio tests.

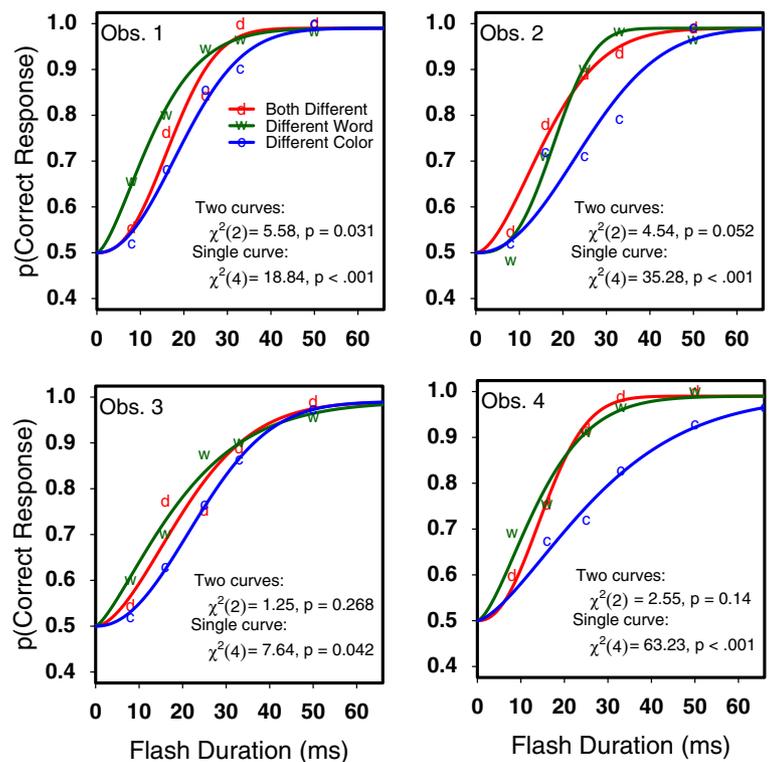


Fig. 3. Color condition's data shown by letters, Weibull fits shown by curves, and likelihood ratio tests (χ^2). Each panel is an observer. Conditions are shown in Table 1. Likelihood ratio tests are between the full model (curves shown) and restricted models (specified in the text). Numbers in parentheses are degrees of freedom for the likelihood ratio tests.

based on size (i.e., case) or color. Such differences could be perceived on the basis of the low-level form of any letter in the string, or even in the absence of identification of any single letter. Instead better performance resulted from a condition requiring identification of a specific (uncued) letter. The results suggest that observers were making a significant proportion of their choices on the basis of higher-level information (such as letter names, letter groupings, word names, or phonology) rather than on the basis of form or color. This conclusion is based both on the fact that different spelling produced better performance than same spelling but different case (or color), and on the fact that for different spelling, case (or color) did not produce additional difference.

Why should lower-level information, such as case or color, be an ineffective source of discriminative information in these tasks? One salient possibility is the nature of the post-mask. For the case conditions the @-signs could have added a good deal of form ‘noise’ so that the only useful information was that encoded to higher levels before the masks arrived. For the color conditions, similarly, the differently colored @-signs could have added color ‘noise’, so that the only useful information was that encoded to higher levels before the mask arrived. Indeed, we noted informally (in pilot testing) that use of black and white @-signs in the color conditions was completely ineffective when the choices differed in color, and performance was at ceiling. Further discussion follows Experiment 2, a study designed to test these ideas by varying the nature of the masks.

3. Experiment 2

Masks are often used to reduce accuracy in perceptual identification studies, under the typical assumption that presentation of the mask terminates persistence of an iconic image (e.g., Sperling, 1960), and thereby stops processing. However, it is also now well established that masks can cause more effects than cessation of processing. In particular, there is good reason to believe that, for tasks like the present one, the features from the mask become combined and confused with the features of the target flash. This idea was explored systematically in short-term visual priming studies by Huber and his colleagues (Huber et al., 2001; Huber, Shiffrin, Lyle, & Quach, 2002; Huber, Shiffrin, Quach, & Lyle, 2002; Wagenmakers et al., 2003; Weidemann, Huber, & Shiffrin, submitted for publication). These researchers developed a model named ROUSE (Huber et al., 2001) that made a large number of correct a priori predictions.

According to the ROUSE model (Huber et al., 2001), features from stimuli are misattributed to other stimuli that are nearby in time and/or space. In the case of the short-term priming paradigm, features from the prime

and the subsequent flashed target word tend to be combined and confused with each other due to their proximity in time and their positions in the display. A logical extension of this assumption to the 2AFC post-masking paradigm posits that the features from the flashed word and the mask might be confused due to their proximity. Of even more relevance to the present studies, the Huber et al. (2001) priming studies used form post-masks, and one of their studies found that case changes between the primes and subsequent target flash produced no performance shift. They suggested that an explanation for this finding might involve the use of higher-level information for decisions, but did not pursue this hypothesis further. In fact, the present studies were motivated in part by the desire to explore this possibility. Our first experiment in this article confirms and considerably extends the findings of Huber et al. (2001). Our second experiment tests the hypothesis that the mask causes the present results by adding noise at a certain low level of information processing. Specifically, Experiment 2 varies the type of mask. Only the case conditions are explored, and three mask types are used: (1) no mask; (2) random pixel noise; (3) the structured pattern masks used in Experiment 1.

If the structured pattern mask selectively disrupts the usefulness of lower-level form features, thus causing the observers to utilize primarily higher-level information, then removal of the mask entirely should allow the observers to return to the use of the normally more effective lower-level form features. If so, the pattern of performance will reverse: Performance should be better in the both-different and different-case conditions relative to that of the different-word condition. The pixel noise mask represents an intermediate level of masking; we included it for generality, but had no clear a priori prediction concerning the results such a mask would produce.

3.1. Methods

3.1.1. Design and materials

The three choice conditions described in Experiment 1 for the case manipulation were crossed with mask type (i.e., @-signs mask, pixel mask, or no mask). All conditions were mixed.

Because elimination of masks produces essentially perfect performance even at the fastest computer monitor refresh rates, we generated psychometric functions by manipulating (log) contrast at fixed flash durations. The use of contrast is probably not a significant change in procedure because in many settings when duration is brief, threshold is a function of both duration and contrast (i.e., Bloch’s Law, Bloch, 1885; Watson, 1986). Pixel masks consisted of a rectangle of pixels that extended just beyond the edges of the target flash. Each pixel subtended approximately 0.04° of visual angle and

had an equal probability of being assigned the background luminance (approximately 1 cd/m²) or maximum luminance. In this experiment, @-sign masks were constructed from the minimum number of @-signs needed to completely cover the target, and the @-signs were set to the maximum luminance. The luminance of the flashed target word was varied against the constant background luminance to produce the desired log contrast, which ranged between -1 and 4.4 . Each observer was run on the same monitor, to control for slight color imbalances across monitors. All monitors displayed stimuli at a 120 Hz refresh rate. The duration of the targets that were not masked was 8.33 ms, while the duration of the targets that were masked with random pixel fields or @-signs was 42 ms. These flash durations were set in a pilot experiment, and thereafter remained fixed.

Six hundred and six one-letter-different word pairs were constructed from medium frequency words (i.e., frequency of 6–100 occurrences per million, as determined by Baayen, Piepenbrock, & van Rijn, 1995) that were 4–6 letters in length. Word pairs were randomly assigned to conditions, position, and target/foil in each session. All stimuli subtended less than 3° of visual angle.

3.1.2. Observers and procedure

Eight Indiana University students, who were not part of Experiment 1, were paid for their participation. These new observers were divided into two groups. Four observers saw all three mask types, and four saw only @-signs masks. The observers that saw all the mask

types participated in 20 sessions each, while the observers that saw only @-signs masks participated in seven sessions each. (The 20 sessions for observers that received all mask types yielded 3600 @-signs mask trials, compared to 3780 @-signs mask trials collected in seven sessions for @-signs mask only observers.) Each session consisted of 18 practice trials followed by 540 experimental trials. Two staircases in each condition for each observer were used to vary the log contrast, one converging to 71% accuracy and the other converging to 84% accuracy. The step size for the staircases was 0.2 units of log contrast. The procedure for each trial was the same as in Experiment 1, except that a centrally-located plus sign (+) was used for fixation.

3.2. Results and discussion

Without a mask, observers used perceptual information to make their decisions (see Fig. 4). Psychometric functions were fit to all conditions for all observers as in Experiment 1. Once a Weibull function was fit, the 75% accuracy threshold was taken from this fitted function. Threshold confidence intervals were derived from threshold fits to 1000 non-parametric bootstrap data sets of the same size as the original data set. The bias-corrected and accelerated 95% confidence intervals were computed from the distribution of thresholds (cf. Efron & Tibshirani, 1993; Wichmann & Hill, 2001). Due to space limitations, these functions are summarized with their 75% thresholds and 95% confidence intervals. (The sets of psychometric functions and likelihood ratio tests

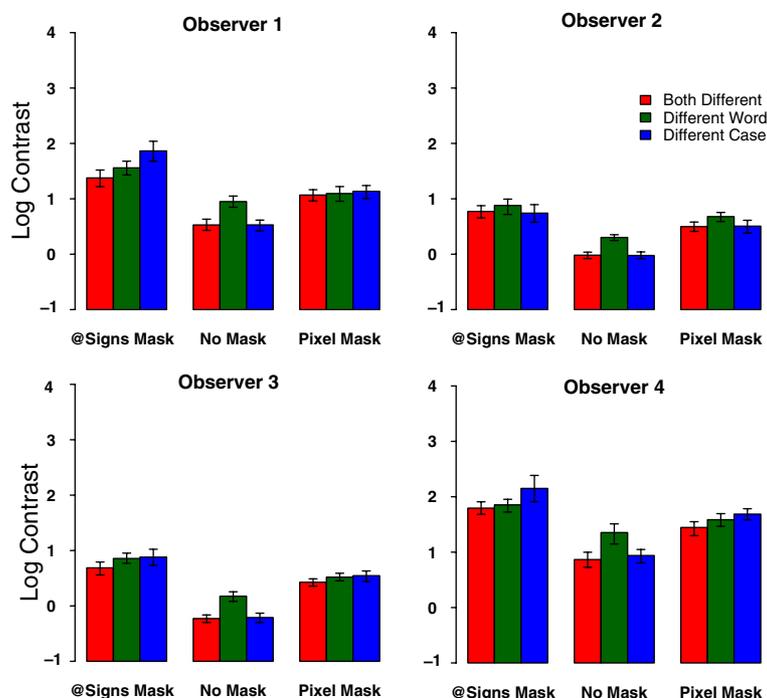


Fig. 4. Thresholds for choice conditions for each type of mask. Error bars around the thresholds are bootstrapped 95% confidence intervals.

such as those from Experiment 1 are available from the authors upon request.)

3.2.1. All mask types

The thresholds for the observers that saw all three mask types are shown in Fig. 4. Comparing the level of performance in the pixel and the @-signs mask conditions shows that the @-signs mask harmed word identification more than the pixel mask. This result indicates that the form features present in the @-signs mask disrupted the availability of the information used to identify the word slightly more than the amorphous pixel masks did. Performance in the no-mask condition is superior to the other conditions, but the size of the advantage is unclear because of the differences in flash duration.

The patterns of performance across the three choice conditions depended on the type of masking. Consider first the no-mask conditions: The both-different and different-case conditions were about equal to each other and superior to the different-word condition, the reverse of the pattern found in Experiment 1 (and to a degree, different from the findings for the @-signs mask condition of Experiment 2). These results suggest that observers in the no-mask condition predominantly use lower-level form information to identify the flashed words because the availability and/or utility of lower-level form features is not selectively disrupted (e.g., unlike in the @-signs mask condition). The results for the pixel field mask were mixed, with no clear and consistent preference for either visual or higher-level information.

For @-signs masks, the results found in Experiment 1 did not consistently replicate. The different-case condition with a pattern mask showed an advantage for observers 1 and 4, as in Experiment 1, but not for observers 2 and 3. It is possible that the mixing of three mask types in Experiment 2 encouraged the utilization of visual form information for decisions even with @-signs masks, in which this strategy is more difficult and/or less effective. That is, there may be a strategic component to the choice of information used to make identification decisions. @-signs masks may generally degrade lower-level form information and induce reliance on higher-level information, but mixing such trials with those in which lower-level form information is useful might tend to produce a uniform strategy of using lower-level form information on all trials.

3.2.2. @-signs mask only

The results of Experiment 1 were not consistently replicated with the observers that saw all three mask types. This result could be due to the influence of the other masking conditions on the @-signs masking condition, or it could be due to other changes in the experimental design between the two experiments. Experiment 2 used contrast instead of flash duration to create the psychometric function and used smaller stimuli than Experiment 1 did. To test the influence of other mask types, four new observers ran in the @-signs masking condition alone. If the reduced preference for high-level information with @-signs masks in the

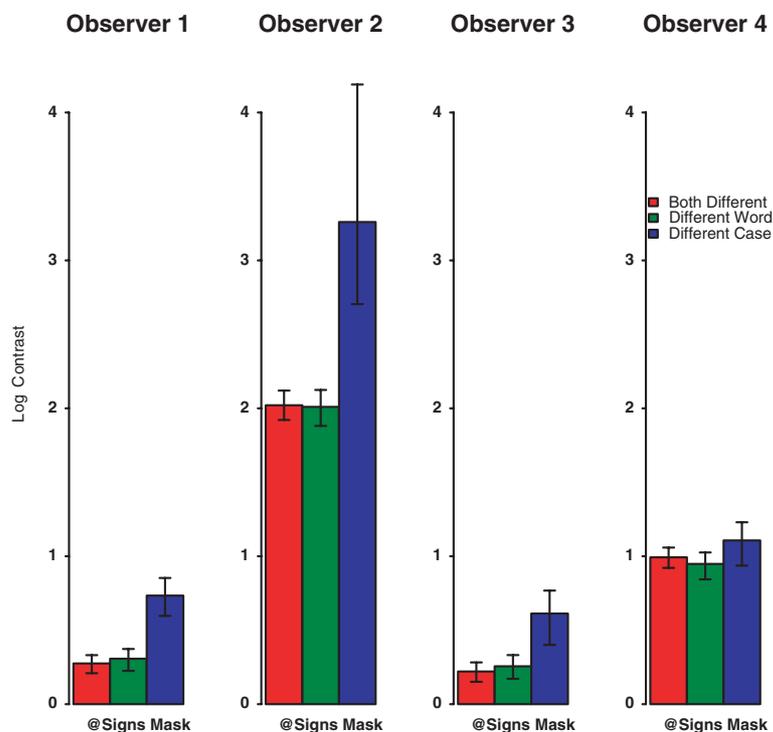


Fig. 5. Thresholds for choice conditions for @-signs masks only. Error bars around the thresholds are bootstrapped 95% confidence intervals.

observers that saw all three mask types versus the observers in Experiment 1 was due to the influence of other masks, the larger preference for high-level information should return in these new observers.

The thresholds for the @-signs only observers are shown in Fig. 5. Thresholds for the different-case condition are larger than the thresholds for the other two conditions, which are nearly identical. This pattern holds for all observers, replicating the results of Experiment 1. Thus, it seems that the failure to replicate the results of Experiment 1 with some observers in Experiment 2 is due to task demands. The variety of masking conditions appears to have changed how observers used information in the @-signs mask condition.

4. General discussion

The pattern of performance across the choice conditions switches quite dramatically as a function of the masks that follow the target flash. Such shifts suggest that the information used to identify briefly flashed words depends on the degree to which masks add ‘noise’ at various levels of information encoding. One interpretation is that the features encoded from the mask are confused with similar features encoded from the flash. If so, the low-level form features would be those confused, degrading the utility of those low-level features to inform decision making. However, before this form mask arrives, processing might proceed to higher levels of information processing (e.g., letter names). This higher-level information is not similar to that in the masks and therefore would not be degraded by the presence of the mask, except to the degree that the mask stops processing. This set of circumstances would naturally lead observers to optimize performance by predominant use of the higher-level encoded features. The use of a mask that does not contain low-level form features similar to that in the target (i.e., pixel masks) seems to attenuate the tendency to use only high-level information for decisions, but does not induce a complete switch to use of low-level information. The result is no clear preference among the choice conditions. The clearest result of course occurs with no mask, because low-level features are no longer degraded, and therefore become the type of information preferred for decision making.

The present study has much in common with research on letter identification. The classic finding from that literature (i.e., the word-superiority effect) is that letters are better identified when flashed as part of a word compared with presentation in non-word, non-letter, or blank contexts (Johnston & McClelland, 1973; Reicher, 1969; Wheeler, 1970). The different-word condition of the present experiments is the same as the word-context condition of the word-superiority studies (at least for studies in which the position of the critical differing

letter is not cued in advance and the choices are words, as in Johnston & McClelland, 1973). Those studies typically use post-masking and should therefore promote use of higher-level visual information, naturally producing a word-superiority effect. The similarity of the paradigms helps to make clear the parallel effects of masks in the two literatures: Removal of masks eliminates the word-superiority effect (Johnston & McClelland, 1973), and it reverses the performance patterns in the present studies. Thus it is reasonable to propose for both sets of paradigms that observers tend to use higher levels of information in the presence of a pattern mask, but lower levels of information when there is no mask. Estes (1975) showed that differences between correct letter identification in word or non-word contexts was due to reduced letter transposition errors for letters presented in word contexts. This result could be taken as evidence that top-down influences on word perception occur after, not during, perception. However, the current results suggest that top-down influences in word perception are concurrent, as visual information is needed in order to activate higher-level units. More formal models are needed to resolve this potential discrepancy.

Repetition blindness also shows results congruent with those found here. Bavelier and Potter (1992) performed a repetition blindness experiment in which the first and second repetitions of a letter were either in the same case or a different case. The effect of repetition blindness was shown by observers not remembering the second instance of the same letter. Observers were specifically instructed to remember the letters presented and their case. However, the effect was similar for a letter presented twice in the same case to that of a letter presented twice in different cases. Further research is needed to determine whether observers would use higher-level visual information when instructed to pay attention to case information.

4.1. Implications for perceptually-driven short-term implicit memory

Implicit memory is revealed when performance in a task that does not require access to any specific past event is nevertheless affected by a recent event. The benefit (or cost) associated with the recent event is referred to as priming. Word identification is one of many tasks that show priming. Tasks like the present ones that use priming are typically termed “perceptual” rather than “conceptual” priming, on the assumption that the recent event affects access to a perceptual or structural representation of the word, as opposed to a conceptual representation of the word.

The 2AFC word identification task has become increasingly important in testing models of short-term implicit memory (e.g., Huber et al., 2001; Ratcliff &

McKoon, 2001). In the short-term priming paradigm, the flashed target item follows the presentation of a priming word after a short delay (e.g., <1 s). The findings from this literature show that priming sometimes facilitates and other times hinders the identification of the briefly flashed target word (e.g., Huber, 2003). The typical 2AFC priming study uses a form post-mask after the target flash, which should promote use of higher levels of information for a decision. According to the ROUSE model (Huber et al., 2001), these higher levels of information include features combined from and confused between the prime and target. Features contributed by the mask are represented in the ROUSE model by a visual noise parameter. To explain the shifts between positive and negative priming, Huber et al. (2001) additionally assume that evidence from features known to have been in primes is discounted to a degree. That is, given that features are confused between primes and targets, it is optimal to discount evidence deriving from features that are known to have been in the prime. Too little discounting produces a bias to respond in accord with the prime, and too much discounting produces a bias against responding in accord with the prime.

Carrying over the present findings to short-term priming paradigms, it is possible to make a prediction: Eliminating the post-mask following the flashed target should make short-term priming case-dependent, assuming there is no mask following the prime either. This prediction has not yet, to our knowledge, been tested. As this discussion illustrates, the present findings point out that a close examination of the test task can provide helpful constraints for models of perceptually-driven implicit memory.

Acknowledgements

The authors wish to thank Adam Reeves and two anonymous reviewers for helpful suggestions on drafts of this article and Jason Gold for advice on using psychophysical procedures.

References

- Baayen, R. H., Piepenbrock, R., & van Rijn, H. (1995). The CELEX lexical database [CD-ROM]. Philadelphia: Linguistic Data Consortium, University of Pennsylvania.
- Bavelier, D., & Potter, M. C. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 134–147.
- Bloch, A. M. (1885). Experiences sur la vision. *Comptes Rendus de la Societe de Biologie*, *37*, 493–495.
- Broadbent, D. E. (1958). *Perception and communication*. Oxford, England: Pergamon Press.
- Broadbent, D. E. (1967). Word-frequency effect and response bias. *Psychological Review*, *74*, 1–15.
- Broadbent, D. E. (1971). *Decision and stress*. London, New York: Academic Press.
- Efron, B., & Tibshirani, R. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.
- Estes, W. K. (1975). The locus of inferential and perceptual processes in letter identification. *Journal of Experimental Psychology: General*, *104*, 122–145.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. Oxford, England: John Wiley.
- Huber, D. E. (2003). The time course of perceptual processing for words and faces. Unpublished manuscript.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002). Mechanisms of source confusion and discounting in short-term priming: 2. Effects of prime similarity and target duration. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *28*, 1120–1136.
- Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, *108*, 149–182.
- Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (2002). Mechanisms of source confusion and discounting in short-term priming: 1. Effects of prime duration and prime recognition. *Memory & Cognition*, *30*, 745–757.
- Johnston, J. C., & McClelland, J. L. (1973). Visual factors in word perception. *Perception & Psychophysics*, *14*, 365–370.
- McGurk, H., & MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, *264*, 746–748.
- McKoon, G., & Ratcliff, R. (1995). How should implicit memory phenomena be modeled? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *21*, 777–784.
- Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly Journal of Experimental Psychology*, *11*, 56–60.
- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, *76*, 165–178.
- Ratcliff, R., & McKoon, G. (2001). A multinomial model for short-term priming in word identification. *Psychological Review*, *108*, 835–846.
- Reicher, G. M. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, *81*, 275–280.
- Schooler, L. J., Shiffrin, R. M., & Raaijmakers, J. G. W. (2001). A Bayesian model for implicit effects in perceptual identification. *Psychological Review*, *108*, 257–272.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*, 1–29.
- Sperling, G., Wurst, S. A., & Lu, Z.-L. (1993). Using repetition detection to define and localize the processes of selective attention. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance xiv: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience—a silver jubilee* (pp. 265–298). Cambridge, MA: MIT Press.
- Treisman, A. (1962). Binocular rivalry and stereoscopic depth perception. *Quarterly Journal of Experimental Psychology*, *14*, 23–37.
- Wagenmakers, E.-J. M., Zeelenberg, R., Huber, D., Raaijmakers, J. G. W., Shiffrin, R. M., & Schooler, L. J. (2003). REMI and ROUSE: Quantitative models for long-term priming in perceptual identification. In J. S. Bowers & C. J. Marsolek (Eds.), *Rethinking implicit memory* (pp. 105–123). London: Oxford University Press.
- Watson, A. B. (1986). Temporal sensitivity. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance*. New York: Wiley.
- Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (submitted for publication). Confusion and compensation in visual perception: effects of spatiotemporal proximity and selective attention.

- Wheeler, D. D. (1970). Processes in word recognition. *Cognitive Psychology*, *1*, 59–85.
- Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, *63*, 1293–1313.
- Zeelenberg, R., Wagenmakers, E.-J. M., & Raaijmakers, J. G. W. (2002). Priming in implicit memory tasks: Prior study causes enhanced discriminability, not only bias. *Journal of Experimental Psychology: General*, *131*, 38–47.