WHAT EXPLAINS A SEMANTIC UNMASKING EFFECT?

by

Alisabeth Ayars

A Thesis Submitted to the Faculty of the

DEPARTMENT OF PSYCHOLOGY

In Partial Fulfillment of the Requirements

For the Degree of

MASTER OF ARTS

In the Graduate College

THE UNIVERSITY OF ARIZONA

2015

STATEMENT BY AUTHOR

The thesis titled *What Explains a Semantic Unmasking Effect?* prepared by *Alisabeth Ayars* has been submitted in partial fulfillment of requirements for a master's degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this thesis are allowable without special permission, provided that an accurate acknowledgement of the source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: *Alisabeth Ayars*

APPROVAL BY THESIS DIRECTOR

This thesis has been approved on the date shown below:

Mary Peterson, Professor of Psychology

Dec. 3rd, 2015
Defense Date

Acknowledgements

I would like to thank Dr. Mary Peterson for her guidance on this research project and the writing of this thesis, Dr. Lee Rylan and Dr. Ken Forster for their useful feedback, and the members of the Visual Perception Lab at the University of Arizona—especially Dr. Joseph Sanguinetti—for helpful discussion throughout this project.

Table of Contents

Abstract

Sanguinetti and Peterson (2013) found that masked words (e.g., “telephone”) followed by a semantically related familiar silhouette (e.g., a silhouette of a telephone) were more likely to be perceived than words followed by an unrelated or novel silhouette, even though the words appeared prior to the silhouettes. This indicates that semantically related items can produce recovery of, or “unmask,” earlier masked items. Sanguinetti and Peterson (2012) interpret their result in a framework proposed by Di Lollo, Enns, & Rensink (2000) which specifies feedback processing to play a role in hypothesis confirmation. According to Di Lollo et al., feedback processing in the visual system constitutes an iterative-loop system, directed at verification of hypotheses about perceptual stimuli. Sanguinetti and Petersonpropose that the familiar silhouettes are able to confirm the hypothesis in the visual system that a semantically related word is present, via reentrant processing. At least two alternative explanations are available for the results. One alternative is that the semantically related silhouettes simply allow for retrieval of the episode of previously seeing the words (which is forgotten upon mask presentation), rather than causally contributing to their phenomenal consciousness. Another alternative is that the semantically related silhouettes allow participants to consciously infer what the word had been from a degraded perception. In Experiment 1, we showed that an attended feature of the words (i.e., their font) is not unmasked along with the semantics of the words—unmasked words appear in no identifiable font. This is evidence against the episodic retrieval hypothesis and supports Sanguinetti & Peterson's original interpretation, since the font of the words would be a component of the original episode of seeing the words. In Experiments 2a and 2b, we show that conscious recognition of the objects in the silhouettes is insufficient for the unmasking effect—reduction in silhouette exposure duration reduces the unmasking effect (Experiment 2a) even though silhouettes are equally recognizable (Experiment 2b). This rules out the inference explanation for the unmasking effect, since on this explanation unmasking would depend on whether the silhouettes are recognized and can

therefore be employed in conscious inference. Independent theoretical contributions of these findings are discussed.

Masking is the reduction in visibility of a target stimulus by an additional stimulus (the mask) presented immediately before or after the target. A variety of kinds of masking have been identified. In random noise masking, the mask elements bear little structural resemblance to the target—the mask, for instance, may consist of random black and white squares. Although some masking is observed under these conditions, masking is much more efficacious if the mask elements approximate the contour properties (e.g., curvature) of the target. This is known as pattern or structure masking. In a third type of masking, metacontrast masking, the target is followed by a mask that does not overlap with the target, but interacts with its contours (e.g., a disk-shaped target followed by a donut-shaped mask that snugly fits around the disk). These forms of masking differ, e.g., in the relationship between target-mask onset asynchrony ("stimulus onset asynchrony"; SOA) and masking efficacy.

It is known that presenting a second mask after the original mask (i.e., a "mask of the mask") allows for full or partial recovery of the target (e.g. Amassian, Cracco, Maccabee, Cracco, Rudell, & Eberle 1993; Breitmeyer, Rudd, & Dunn, 1981; Briscoe, Dember, & Warm, 1983; Robinson, 1966; Tenkink, 1983; Ogmen, Breitmeyer, Todd, & Mardon, 2006). This phenomenon is known as 'target recovery," or "unmasking." Unmasking has been shown for both metacontrast masking (Ogmen et al. 2006) and pattern masking. The presentation of a homogeneous stimulus after a pattern mask reduces the deleterious effects of the pattern mask on target perception (Woodman, Di-Yi, Chun, & Schall, 2004).

Researchers investigating unmasking have manipulated the physical properties of the second mask, but not the semantic relation between the target and second mask. Recently, Sanguinetti and Peterson (2013) demonstrated a semantic unmasking effect, in which the presentation of silhouettes semantically related to masked target words resulted in target recovery. Sanguinetti and Peterson (2013) used pattern masks as the first mask, but rather than present a homogeneous stimulus as a second mask, they presented participants with words semantically related to the target words.

Participants were first presented with words that named animals, plants, or objects for 50 ms (see *Figure 1*). The words were forward masked with hashtags and backward masked with random letter strings that were the same length as the word. The words were thus usually unperceived. Then, a silhouette was briefly presented. There were both Familiar and Novel silhouettes. For Familiar silhouettes, the silhouette depicted an object that is recognizable and nameable, and that participants likely would have encountered outside the laboratory. For Novel silhouettes, the silhouette did not depict a recognizable, nameable object. There were two kinds of Novel Silhouettes: Control and Experimental. Experimental Novel silhouettes suggested portions of familiar objects on their outer borders that participants did not consciously perceive (e.g, the profiles of seahorses in the "Experimental Match" condition in *Figure 1*) because they were perceived as part of the background rather than the figure. Control Novel silhouettes depicted portions of unnamable, unrecognizable objects on their groundsides. For the Familiar silhouettes and the Experimental Novel silhouettes (suggesting familiar objects on their groundsides), half the words "matched" the silhouettes, meaning that they named the familiar object in the silhouette (or, in the case of Experimental Novel silhouettes, the objects suggested on the outer borders). The other half of the Familiar and Experimental Novel silhouettes "mismatched" the word, meaning the word named a different object than the familiar object(s) in the display. The silhouettes were presented for 175ms and pattern masked. There were 20 trials for each condition (Familiar Match, Familiar Mismatch, Experimental Match, Experimental

Mismatch, and Control), and participants received all trial types.

Participants' main task was to categorize the silhouettes as “Novel” or “Familiar.” They were asked, after each trial, whether they saw a word on that trial. They were told to only report seeing a word if they had a clear perception of it (i.e., if the word was phenomenally conscious). Participants were also asked, if they saw a word, whether it matched the silhouette, in order to assess whether participants perceived the words accurately.

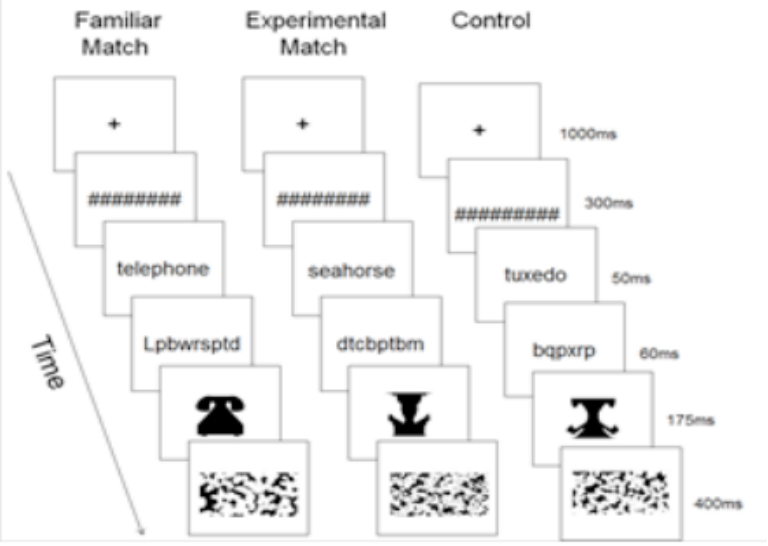

Figure 1 Trial structure in Sanguinetti & Peterson (2013). The trials depicted are from the Familiar Match condition, in which the word (“telephone”) named the silhouette, the Experimental Match condition, in which the word (“seahorse”) named the objects depicted on the outer borders of the silhouettes (the white regions in the silhouette suggest profiles of the upper bodies of seahorses facing inward), and the Control condition. There were also mismatch conditions for the Familiar and Experimental conditions. An example of a word utilized in the Familiar Mismatch condition along with the telephone silhouette is “chocolate.”

Participants were more likely to report seeing words when the silhouette was Familiar and matched the word than in all other conditions (*Figure 2*). This indicates that the semantically related Familiar silhouettes allowed for the recovery of the masked words that named the silhouettes. One might worry that participants simply guessed that they saw the word that named the silhouette whenever a familiar silhouette appeared. However, this concern is preempted by the fact that participants did not see more words in the Familiar Mismatch condition, even though the silhouettes in this condition were also Familiar. Furthermore, when participants did report seeing words in the

Familiar Mismatch condition, they usually accurately reported that the words did NOT match the silhouette (see *Figure 2*). There is therefore no evidence that participants are biased simply to report that they saw words matching the objects in the Familiar silhouettes.

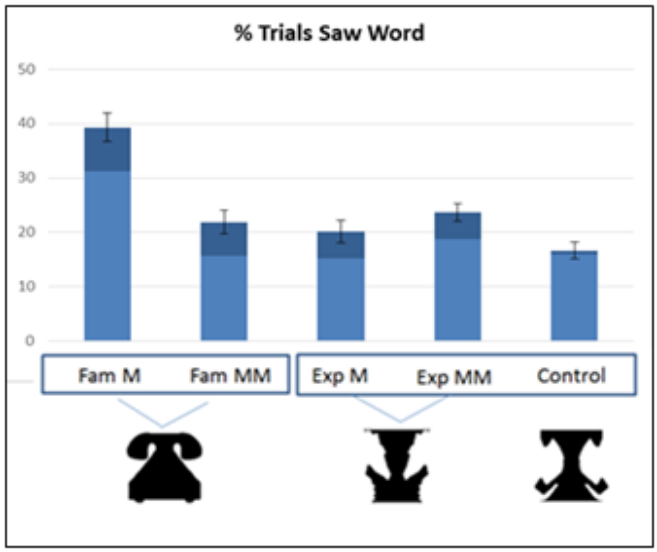

Figure 2 Results of Sanguinetti & Peterson (2013). FamM = Familiar Mismatch, FamMM = Familiar Mismatch, ExpM = Experimental Match, ExpMM = Experimental Mismatch. Word perception was significantly higher in the Familiar Mismatch condition than all other conditions. The light blue areas of the bars portray the percent of trials in which participants were accurate in whether the word matched the silhouette. Accuracy was generally high and did not fall below 70% in any condition.

The Experimental novel matching silhouettes, which depicted familiar objects on their outer borders, did not allow for unmasking. This is somewhat surprising, in light of results demonstrating that potential objects ultimately assigned to ground are nevertheless semantically processed (Peterson, Cacciamani, Mojica, & Sanguinetti, 2012; Sanguinetti, Allen, & Peterson, 2013). It seems that semantic activation matching that of the word is not sufficient to facilitate word perception; the semantically related region must be assigned object status as well.

This "semantic unmasking effect" is an exciting finding. It cannot be explained straightforwardly by appealing to the ability of the semantically related silhouettes to mask the masks of the words (i.e., the random letter strings), which is the standard interpretation of unmasking effects (e.g., Ogmen, Breitmeyer, Todd, & Mardon, 2006). The silhouettes do not structurally resemble the

contours of the random letter masks and thus are unlikely candidates for being structure or pattern masks for the random letter strings. Furthermore, and crucially, the unmasking effect depends on the semantic relation between the target word and the subsequent silhouette. The present research attempts to discern competing explanations for the effect. We also advance research on the semantic unmasking phenomenon in a few ways. First, we show that features of words (namely, their font) are not unmasked. Furthermore, we show that conscious identification of the silhouettes is insufficient for unmasking, and the unmasking effect is lessened if silhouette processing is reduced.

Iterative Processing and Hypothesis Confirmation

Sanguinetti and Peterson (2013) draw on theory developed in the literature on object substitution masking to explain the semantic unmasking phenomenon. Enns & Di Lollo (1997) discovered that four dots that surround a target, but do not touch the target or interact with any of its contours, can mask the target (see *Figure 3*). This is known as *four-dot masking* or *object substitution masking*. In this form of masking, a target (such as a circle) onsets with four dots surrounding the target. Then, the target is offset, but the dots remain. To the participant, it appears as if the target has been "substituted" for the four-dots too quickly for target identification to occur. Masking does not occur when target and mask offset together—in this scenario, the target is easily seen.

Certain features of four-dot masking make it difficult to explain using standard theories of masking, such as the dual-channel approach or temporal integration (see Enns & Di Lollo, 2000, for a

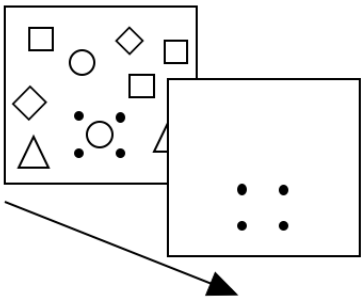

Figure 3 Object substitution masking

thorough presentation of these arguments). In four-dot masking, the mask does not interact with the target's contours—the spatial proximity of the dots to the target in fact matter little for object substitution masking, unlike normal masking (e.g., metacontrast masking). The efficacy of four-dot masking requires that attentional resources be distributed. This is why the target circle in *Figure 3* appears with many other shapes and is not centrally located. Finally, quite strikingly, masking can still occur if, strictly, the dots offset with the target but *apparently* do not, via an apparent motion illusion. Lleras & Moore (2003) found that targets were still masked by four dots when the four dots apparently "moved" from one part of the display to another once the target offset. (In actuality, a new set four dots were simply onset after the original four dots offset, creating the illusion that the original dots had moved to the location of the later four dots).

In order to account for these feature of object substitution masking, Enns & Di Lollo (2000) and Di Lollo, Enns, and Rensink (2000) propose that object substitution masking occurs because the presence of the target fails to be confirmed via feedback processing when the four-dots linger after target offset. According to the authors, feedback processing constitutes an iterative-loop system, aimed at verification of hypotheses about perceptual stimuli. The job of this system is to establish the most plausible interpretation of the incoming stimulus. This is accomplished by continuously comparing, via iterative processing loops, the neural activity *predicted* in lower visual areas on a given hypothesis about a stimulus (such as the presence of a circle) to the *actual* activity in those areas. A match between predicted and actual activity acts as verification of the hypothesis of the stimulus, rendering it conscious after sufficient iterations.

Applied to object substitution masking, upon the presentation of the target and mask, neurons coding information about the stimulus at that location register the presence of both the target and mask. Therefore, the visual system hypothesizes the presence of the target-mask combination at that spatial location. However, when the target-mask hypothesis is "re-checked" (via reentrent processing), only

the presence of the mask is confirmed, since the target offsets so quickly. The visual system thus favors the interpretation that the stimulus at that location had always been the four dots only, and the original target+mask hypothesis is discarded. Masking does not occur when the four-dots offset with the target because in that case, the only stimulus that is confirmed is the target+mask combination (i.e., the visual system never detects the presence of four-dots alone). Masking does not occur when the dots onset after the mask because in that case the four-dots is treated as a separate object than the target.

This framework accounts for the distinguishing features of object four-dot masking. For instance, attentional resources must be distributed because if attention is focused on the target, sufficient iterative loops of hypothesis confirmation can occur for the presence of the target to be confirmed. Apparent motion between four-dots surrounding the target and four-dots onsetting in a different location is sufficient for masking because in such a case the two sets of dots are interpreted as being the same object.

Sanguinetti and Peterson build on the framework of iterative processing and hypothesis confirmation of perceptual stimuli to explain the unmasking effect. Although they used pattern masking and not OSM masking, similar principles may apply. They propose that upon presentation of the word, a hypothesis is formed by the visual system that the word is present. The hypothesis includes semantic information as well as information about word form. When the hypothesis is re-checked via reentrant processing, it is disconfirmed because the word has been replaced by the random letter mask, rendering the word invisible. However, crucially, Sanguinetti and Peterson argue that the hypothesis of the word is still active even after this initial disconfirmation. The visual system continues to re-check the consistency of the data (i.e., lower-level neural activation) with the word hypothesis for some amount of time. When the silhouette appears, it activates a hypothesis with the same semantic content. The activation of an hypothesis with the same semantic content as the hypothesis for the word causes the hypothesis regarding the word to exceed a threshold required for perceiving the word. If the silhouette

is Novel or semantically unrelated to the word, it does not activate a hypothesis with the same content as the original word, explaining why there is no unmasking in the conditions using Novel silhouettes and in the Familiar Mismatch condition.

If Sanguinetti & Peterson's explanation for the unmasking effect is correct, it would indicate that confirmation of a perceptual stimulus can occur via a consistency of *semantic* information as opposed to just featural consistency in a particular spatial location. In standard four-dot masking, the presence of the four-dot object is confirmed because the four dots remain after target offset, and were also present along with the target. In Sanguinetti & Peterson, there is no featural similarity between the words and the matching silhouettes. Nevertheless, the incoming semantic activation from the silhouette is treated as evidence of the presence of the word, if Sanguinetti and Peterson's interpretation is correct. This would be a significant and exciting extension of the theoretical framework.

However, at least two alternative explanations for the results are available. The first involves memory. It may be that participants have a conscious perception of the word on most trials but this experience is immediately forgotten upon presentation of the random letter mask. When a semantically related familiar silhouette appeared, it triggers participants' memory of the experience of seeing the word. This process has been argued to occur in repetition priming (e.g., Logan, 1990; Neill, 1997). On this account of repetition priming, the first word presented in a repetition sequence is encoded in episodic memory, and the second presentation causes retrieval of the first word (Logan, 1990). One piece of evidence for masked primes constituting episodes is that they are effective in a “mere exposure” paradigm—brief, masked exposures can lead subjects to prefer that prime to novel one, despite chance performance at prime recognition (Kunst-Wilson & Zajonc, 1980; Seamon, Brody, & Kauff, 1983; Whittlesea & Price, 2001).

According to the memory explanation of Sanguinetti and Peterson’s results, the silhouette does not facilitate the word becoming conscious—rather, it simply reactivates the memory trace of an earlier

experience of seeing the word. This makes this explanation different than Sanguinetti and Peterson's, which specifies that the presence of the Familiar silhouettes is causally important to the semantically related words reaching consciousness at all. In other words, on Sanguinetti & Peterson's account, the Familiar silhouettes cause the words to be perceived; on the memory account, the words are more-or-less always perceived, but are only recalled in a manner that makes them available for report if followed by a Familiar matching silhouette.

If the memory interpretation is correct, this would not be without theoretical import. For instance, it would suggest that fully conscious stimuli can be nearly immediately forgotten. This is a conclusion pushed by Chen & Wyble (2015), who found that participants were unable to report on an object feature that they had presumably needed to be conscious of just milliseconds before in order to perform a task. However, crucial to Chen & Wyble's effect was that participants were surprised by the request to report on the object attribute; Chen & Wyble achieved this surprise by unexpectedly asking participants to report the feature after many trials performing a different task. We think therefore it is unlikely that participants had an experience of the word that was immediately forgotten on each trial because their task set was to report whether they saw a word and is thus disanalogous to Chen & Wyble's paradigm. Nevertheless, the memory interpretation warrants consideration.

A second alternative explanation involves conscious, deliberate inference. According to this explanation, participants may have had a conscious perception of the word that was degraded or incomplete due to the presence of the mask. However, if some identifiable letters of the word survived, the familiar matching silhouette may have allowed for an inference about the identity of the word. For example, if participants had a degraded perception equivalent to "t_lep_one," they might infer that word had been "telephone" upon the presentation of the telephone silhouette. According to this explanation, reports of seeing a word were higher in the *Familiar Match* condition because it is only in this condition that the familiar objects matched the degraded word perception and could thereby

facilitate inferences about the words.

The Present Project

The current project aims to tease apart the competing explanations for Sanguinetti & Peterson's results as well as further investigate the dynamics of the unmasking effect. The episodic retrieval hypothesis posits that Familiar matching silhouettes trigger a retrieval of the episode of seeing the word. If this is right, the word should be recalled along with the features encoded during the episode of seeing the word. In Experiment 1, we varied the font of the words and asked participants to report on the font of each word, as well as whether they saw a word on each trial. Since font report was in their task set, participants should be expected to attend to and encode the font of the words if the words were conscious upon presentation. Thus, they should be able to report on the font of the word if the familiar silhouettes simply facilitate the retrieval of the experiential episode of seeing the word.

Experiment 1 also allows us to test the independently theoretically interesting question of whether features of the word are unmasked as well. If it turned out that participants are better able to report on the font of the words in the Familiar Match condition, this would suggest either that the episodic memory hypothesis is correct, or that hypothesis confirmation operates over detailed representations that preserve many features of the original stimulus (such as its font). If participants are not able to report on the font of unmasked words (which is what we found), this suggests that hypothesis confirmation operates (or at least can operate) over relatively coarse-grained representations or "abstractions" of the original stimulus.

In Experiments 2a and 2b, we explored whether recognition of the silhouettes is sufficient for the unmasking effect. According to the inference explanation, participants infer the presence of the words based on the semantically related silhouettes that follow. The "unmasking" effect is thus produced by conscious recognition of the silhouettes on this explanation. We predicted that the unmasking effect would be reduced when the silhouettes were exposed for 70ms instead of 175ms,

despite silhouette identification being equally good at 70ms. In accordance with this prediction, we found that word perception decreased in the *Familiar Match* condition when the silhouettes were exposed for 70ms rather than 175ms (Experiment 2a), despite that overall the silhouettes were equally recognizable at this exposure duration (Experiment 2b).

Reducing silhouette exposure also provides a window into the extent to which the silhouettes must be processed in order for the unmasking effect to occur. On the hypothesis confirmation account, less silhouette processing entails that less iterations confirm the presence of the semantic item, and thus would be expected to reduce the unmasking effect. The Experimental Novel silhouettes, containing familiar objects on the outer borders of the novel central objects, did not allow for unmasking in Sanguinetti and Peterson's (2013) experiment. This is perhaps because the outer shapes were not confirmed as objects in the display; only the central region was assigned object status. If reduced silhouette processing entails that even the central region is not confirmed to the same extent as an object in the display, then the Familiar silhouettes may be less able to facilitate word perception.

Experiment 1

Experiment 1 aimed to untangle the episodic retrieval explanation from the hypothesis confirmation account by requiring participants to attend to and report on a feature of the word—that is, whether it was in italic or regular font. The episodic retrieval hypothesis specifies that Familiar matching silhouettes trigger a retrieval of the episode of seeing the related word. If this is right, the word should be recalled along with the features encoded during the episode of seeing the word (especially since there would only be a delay of milliseconds between exposure to the word and the report. If the episodic memory explanation is correct, then the extra words recalled in the Familiar Match condition should be recalled along with their font, given that participants are asked to attend to the font of the words and therefore would be expected to see and encode this feature of the words. Accuracy at reporting the font of the words should therefore be higher in the Familiar Match condition.

We utilized the same paradigm as Sanguinetti and Peterson (2013), except 50% of the words were in italic font, and 50% were in normal font. Participants were asked on each trial whether the word was in italic or regular font, and whether the word matched the silhouette.

Participants

72 undergraduate students from the University of Arizona participated in this study after giving informed consent. Participants completed the experiment to fulfill a requirement for their introductory psychology course. Participants all reported normal or corrected-to-normal visual acuity.

22 participants reported seeing objects on the groundside of the Experimental Novel silhouettes. They were excluded from the main analyses, although their data was analyzed separately. This left 50 participants in the main analysis.

Stimuli and Apparatus

A personal computer with a 21 inch Sony CRT monitor with was used to display the stimuli and record responses. The presentation software was DMDX (Forster & Forster, 2003). Participants viewed the monitor from a distance of 96 cm using a vertically adjustable chin rest to maintain head position. Participants responded using a left-right button box. To initiate the start of each trial, the participants pressed a foot pedal.

The critical visual stimuli were symmetrical silhouettes that portrayed either familiar or novel shapes. (See *Figure 4*). These stimuli are the stimuli used in Sanguinetti and Peterson (2013) and were developed by Peterson and Kim (2001), Peterson and Skow (2008), and Trujillo, Allen, Schnyer, & Peterson, 2010. Familiar shapes were plants, objects, or animals that participants presumably would have encountered outside the laboratory. The Novel silhouettes bore little resemblance to objects that participants would have previously encountered. In all the displays, the center black region (the silhouette) was symmetric and smaller than and enclosed by the surrounding white space, entailing that it is normally seen as the figure. 60 Novel silhouettes and 40 Familiar silhouettes were used. There

were two types of Novel silhouettes: 40 Experimental Novel silhouettes and 20 Control Novel

Figure 4 Examples of stimuli used in Sanguinetti & Peterson (2013) and in the present project. The Familiar silhouette depicts a telephone. The Experimental Novel silhouette suggests the profiles of the upper body of seahorses facing inward on the outside of its borders. The Control Novel silhouette depicts unique objects on the outside of its borders.

silhouettes. (Twice as many Experimental Novel as Control silhouettes were used because Experimental Novel silhouettes were divided into 2 conditions, Match and Mismatch). Again, although we did not expect to find an effect of Novel silhouette type, we used both types to mirror the conditions in Sanguinetti & Peterson's original experiment. The left and right vertical borders of the Experimental Novel silhouettes suggested a portion of a familiar shape in the outer regions. The vertical borders of the Control Novel silhouettes suggested a novel object in the outer region. The Experimental Novel and Control Novel silhouettes were equated on the following features: contour length, luminance, horizontal span, spatial frequency, symmetry, enclosure, and convexity (Trujillo et al., 2010). Familiar silhouettes were matched as closely as possible to Novel silhouettes; see the appendix of Trujillo et al. for details. Participants were expected to perceive the inside region as figure even for Experimental Novel silhouettes because there are stronger cues favoring the inside shape as figure, such as small area, symmetry, and enclosure. Silhouettes were matched in height across types (4.7°H X 1.8°–9.4° W).

Design and Procedure

Instructions were presented to the participants on the monitor as well as being read aloud. The instructions described the trial structure with two qualifications: participants were told that most but not all trials would include words (with non-word trials containing random letter strings in place of words),

and that there were only two types of silhouettes (Familiar and Novel). Thus, participants were not informed that some novel silhouettes contained familiar objects on the outer borders, nor that some silhouettes would match the words, nor that words were presented on all trials. Participants were told that their task was, for each trial, to report 1) whether the word or random letter string was italic or regular font, 2) whether they had seen a word on that trial and, if they answered "yes" to question 2, 3) whether the word matched the silhouette.

There were 100 trials in total for each participant. Each trial began with a fixation cross at center. The participant pressed the foot pedal to begin the trial. After a one second delay, a forward mask consisting of hashtags (#####) was presented for 500*ms. The string of hashtags was the same length as the subsequent word. Then, a word was presented for 83*[1]ms. Words were matched for length and frequency and named various objects, animals, and plants. Half of the words were in italic font, and half were in regular font. Two scripts were created with opposite words in italic font to ensure that any effect of particular words was controlled. The word was backward masked by a string of random letters the same length as the word, presented for 100* ms. Then, a silhouette was presented for 283* ms. The silhouettes were followed by a pattern mask, presented for 666*ms (See *Figure 5* for a visual depiction of the trial structure). A response screen then appeared, prompting the participant to report whether the word or random letter string was in italic or regular font. After making a response, the participant was prompted by a different response screen to indicate whether they saw a word on that trial, or just random letters. If the participant responded "word," another response screen appeared in which the participant was asked to indicate whether the word matched the silhouette. Each response screen was maintained until the participant made a response using the button box or until the trial

1 The presentation duration of 83ms was an error. Because the video mode refresh rate was unknowingly set to 60hz rather than 100z, each display was presented for 167% of the planned duration.

timed-out (which occurred 3000ms after the onset of each response screen).

There were five within-subject conditions, with 20 trials per condition. These were identical to Sanguinetti & Peterson's conditions. In the *Familiar Match* condition, the silhouette was familiar and depicted the object named by the word (N = 20 trials). In the *Familiar Mismatch* condition, the silhouette was familiar but did not match the word (N = 20 trials). In the *Experimental Match* condition, the familiar object suggested on the groundside of the Experimental Novel silhouette was named by the word (N = 20 trials), and in the *Experimental Mismatch* condition, the groundside regions depicted a different object than the one named by the word (N = 20 trials). In the *Control* condition, no familiar objects were suggested on the groundsides of the silhouettes and thus there could not be a semantic relationship between the words and silhouettes (N = 20) trials. Trials were randomized for each participant.

Before the experiment began, participants completed a set of 10 practice trials that utilized only Familiar and Control Novel silhouettes (the Familiar silhouettes mismatched the word for all practice

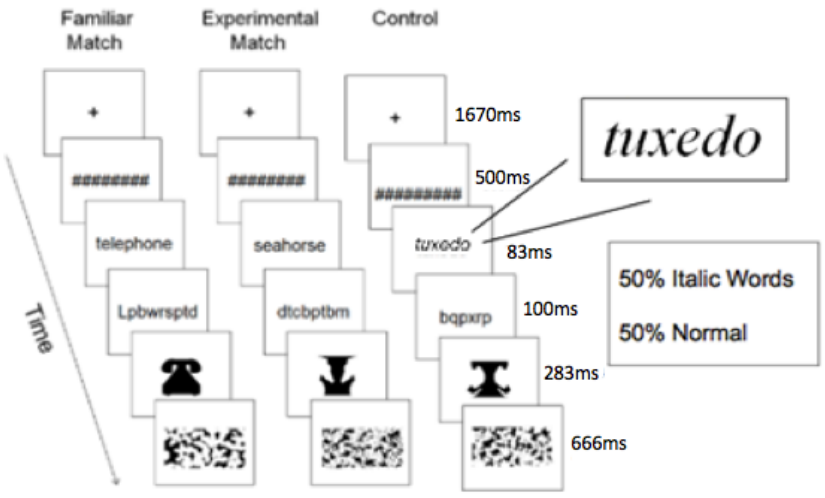

Figure 5 Trial structure of experiment 1a.

trials). An experimenter remained in the room to answer questions during these practice trials. Participants were not given any feedback.

Following the experimental trials and before debriefing, we asked subjects a series of post-

experiment questions to determine whether they had seen any of the objects suggested on the outsides of the silhouettes, consistent with Sanguinetti and Peterson (2013). Participants were first asked whether they noticed anything unusual about any of the silhouettes, and if they did not report noticing groundside objects, the participant was shown a sample Experimental Novel silhouette and the groundside objects (termed "hidden objects" for participants) were pointed out. Then, participants were asked if they had seen any "hidden objects" during the experiment. The majority of participants (50/72) did not report seeing groundside objects. The data of participants who saw the groundside objects were excluded from analysis as per Sanguinetti and Peterson (2013), since seeing the outside objects as figure could potentially affect responses. Since such a large percentage of participants were aware of the objects suggested on the groundside of the Experimental silhouettes (31%), their data were analyzed separately. We present the results of that analysis after the analyses of the participants who were unaware of the objects suggested in the grounds, with respect to each test.

Results

Font Report Accuracy

Figure 6 shows the mean and standard error bars for the percent of trials in which participants accurately reported the font of the word in each condition. No difference in font accuracy was found

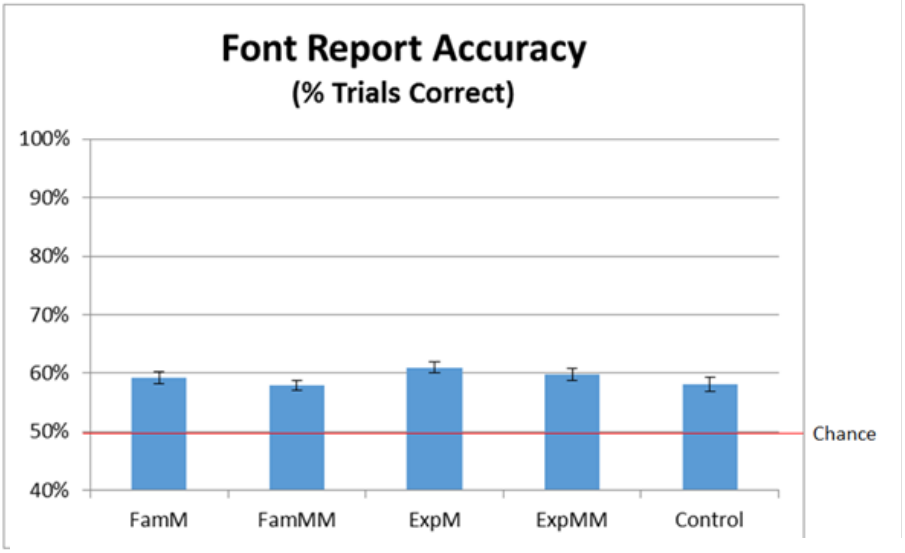

Figure 6 Font report accuracy, experiment 1.

based on which set of words were italic, so these groups were collapsed. Overall, font report accuracy was low. Chance is 50%, and accuracy did not exceed 61% in any condition. Nevertheless, font accuracy was greater than chance for all conditions ($p < .001$). A one-way ANOVA with the factor of condition (*Familiar Match*, *Familiar Mismatch*, *Experimental Novel Match*, *Experimental Novel Mismatch*, and *Control Novel*) revealed no significant effect of condition for font accuracy, $F(4,49) = 1.20$, $p = .31$.

Overall mean font report accuracy was not different for seers (i.e., participants who reported seeing objects on the groundsides of the silhouettes when asked in a post-experiment debriefing session) ($M = 60\%$) than for non-seers. However, a one-way ANOVA revealed an effect of condition on accuracy. Surprisingly, mean font accuracy in the Experimental Mismatch condition, although still relatively low ($M = 65\%$), was higher than all other conditions for seers, $F=3.46$, $p = .01$. This may be a statistical artifact, due to the low sample size for the seer group (N = 22).

Word Perception

Figure 7 shows the mean and standard error bars for the percent of trials in which participants reported seeing a word. A one-way ANOVA with the factor of condition (*Familiar Match*, *Familiar Mismatch*, *Experimental Novel Match*, *Experimental Novel Mismatch*, and *Control Novel*) revealed a significant effect of condition, $F(4,49) = 36.31$, $p< .001$. A post-hoc LSD test adjusted for multiple comparisons found that word perception in the *Familiar Match* condition (M = 68.9%, SE =2 %) was higher than word perception in all other conditions, $p < .001$, as expected.

Word perception was overall higher for seers (M = 63%) than for non-seers (M = 52%). Mean word perception in the Familiar Match condition for seers was 81%, and was higher than all other conditions, $F(4,21) = 25.87$, $p < .001$, indicating that results for seers replicated Sanguinetti and Peterson's unmasking effect, as did non-seers. Despite reporting seeing objects on the groundsides of the Experimental Novel silhouettes, word perception was not higher in the Experimental Match condition (M = 67%) than for the other conditions for the seer group.

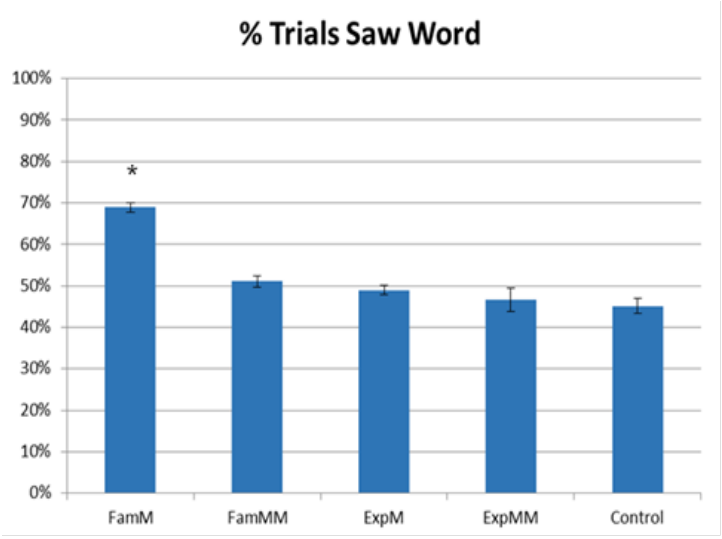

Figure 7 Word perception, experiment 1. Word perception was higher in the Familiar Match condition than all other conditions, replicating Sanguinetti & Peterson's original finding.

Match Judgment Accuracy

On each trial that participants saw a word, they were asked to indicate whether it matched the silhouette. The correct answer to whether the word matched the silhouette is "no" in all conditions except the *Familiar Match* condition. (The correct answer is "no" in the *Experimental Match* condition because participants were asked to make the judgment about the central white silhouette). Overall, participants were accurate in this judgment; accuracy did not fall below 77% correct in any condition. For the purposes of the ANOVA, we analyzed the proportion of "yes" responses rather than the proportion correct. Because "yes" responses appeared higher in both match conditions (Experimental Match and Experimental Match) compared to the mismatch conditions, we conducted a 2-way

ANOVA with the factors of "silhouette type" (Experimental vs. Familiar) and "match" (Match vs. Mismatch). The ANOVA indeed revealed a significant main effect for the "silhouette type," $F(4, 49) = p < .001$) and "match" factor ($p < .001$), as well as an interaction ($p < .001$). The interaction suggested that the difference between the *Familiar Match* and mismatch conditions (difference = 62.96%) was greater than the difference between the match and mismatch conditions for the experimental silhouettes (difference = 10.53%). This interaction was expected, since "yes" is the correct answer only in the *Familiar Match* condition. The difference between the Experimental Match and Experimental Mismatch conditions (difference = 10.53%) was also significant, $t=2.74$, $p < .01$). This suggests that the hidden shapes on the groundside occasionally led participants to mistakenly report that the word matched the silhouette, despite not perceiving the shapes.

Discussion

Experiment 1 replicated the unmasking phenomenon found by Sanguinetti and Peterson (2013). Word perception was higher in the *Familiar Match* condition than all other conditions. Word perception was not higher in the Experimental Match condition than the other conditions, suggesting that the familiar objects in the silhouettes must have figure status to facilitate word perception. Sanguinetti and Peterson (2013) also found no effect of the Experimental Novel silhouettes on word perception. The present experiment suggests that increasing the exposure duration of the silhouettes to 283ms does not change this.

Crucially, font report accuracy was no higher than in *Familiar Match* condition than the other conditions, and only slightly exceeded chance. This does not support the episodic retrieval explanation for the unmasking effect. If the Familiar silhouettes were simply triggering a memory of the experience of seeing a word, then the font of the word—since it would be part of the episodic memory of seeing the word—should be reportable along with the semantics of the word when the word is "unmasked" in the *Familiar Match* condition.

The exposure durations for Experiment 1 were longer than in Sanguinetti and Peterson (2013), due to a mistake in setting the monitor refresh rate. We decided to include the experiment despite the increased exposure durations for a few reasons. For one, although participants ended up seeing many more words in the present experiment than Sanguinetti and Peterson (2013) as would be expected with a longer exposure duration for the words, we still obtained the unmasking effect in the Familiar Match condition. Secondly, the longer exposure duration of the words helps to assuage a possible concern about the experimental design—that the short exposure duration would not have in principle allowed for perception of the font of the word, because there was insufficient time for the font to be seen. If the original episode of the words did not include the font, then the font would not be expected to be reportable upon retrieval of the episodes, rendering the present experiment an ineffective test for ruling out the episodic memory hypothesis. The longer exposure duration of the words mitigates (although does not entirely preempt) this concern.

We found a large proportion of seers in this experiment. This is likely due to the long exposure duration of the silhouettes (283ms, compared to 170ms in Sanguinetti and Peterson's experiment). Word perception overall was also higher in this experiment than in Sanguinetti & Peterson's experiment. Participants saw about 50% of the words in all the conditions except Familiar Match, in which they saw about 70%. We think this is also explained by the longer exposure duration of the words. Other factors may have contributed as well, such as the different task of the experiments. In the present experiment, participants were first asked to report the font of the word on each trial. In Sanguinetti & Peterson's original experiment, participants were first asked to categorize the silhouette. Participants' primary task was thus different across the two experiments—in the present experiment, it was directed at the word. The increased attention to the word in the present experiment may have facilitated greater overall word perception.

The fact that font reports were close to chance in all conditions despite the 83ms word exposure

duration may suggest that font was not particularly visible. It is important to address this concern, since it is relevant to whether this experiment is a fair test of the episodic retrieval hypothesis. As noted earlier, if the exposure duration of the words was too short for the words to be visible even on the assumption that masking may not eliminate the visibility of briefly presented stimuli, then proponents of the episodic retrieval hypothesis might contend that the font was not retrieved with the word semantics in the Familiar Match condition simply because it was not seen (and therefore not encoded).

It is important to remember, however, that the memory retrieval explanation predicts that the font of most words will be un-reportable, since most words (along with their font) are immediately forgotten due to the presence of the random letter mask. In other words, proponents of the episodic memory explanation cannot explain away the results by contending that low font report accuracy indicates that the font of the words were not seen. It may be that most words (and their font) are *seen* but *forgotten*, precisely as the episodic memory hypothesis predicts.

Furthermore, a few features of the experimental design facilitate the font of the words being visible (under the assumption that masked stimuli presented at that duration are capable of being seen in the first place). Participants' primary task was to report on the font of the word, as it was the first question they were asked on each trial. Thus, participants would be expected to attend to the font of the word on each trial. Finally, as noted, due to a mistake in the refresh rate setting on the video mode, the words were presented for a longer time than in Sanguinetti and Peterson (2013)—83ms rather than 50ms. The lengthened exposure duration would facilitate the font being perceived, compared to a shorter exposure. Therefore, under the assumption (made by proponents of the episodic retrieval hypothesis) that masked primes are in fact visible but merely forgotten, it seems that conditions in the present experiment were optimal for font perception.

There is reason to be concerned that despite the optimal conditions, however, the font was particularly difficult to see. For instance, even when participants reporting seeing words in the

conditions other than Familiar Match (i.e., when they were able to see the words despite the mask, and without the aid of the Familiar silhouettes), they often were not able to report the font. (If participants were able to see the font of the words every time they reported seeing a word, one would expect font accuracy to be much higher—around 75%).

We accept this concern as valid. The fact that participants saw many words that they were unable to report the font of in all conditions was surprising—we expected that in most cases in which participants saw words, they would also see the font of those words. However, we think the results of the present experiment still counts as some evidence against the episodic retrieval explanation of Sanguinetti and Peterson's unmasking effect because even though font report accuracy was unexpectedly low, *no* increase in accuracy was observed in the Familiar Match condition. Even if font was generally difficult to see, it was visible at least sometimes—font report accuracy was greater than chance.

One can calculate that the font was visible on average in 20% of trials. This is calculated by solving the following equation, in which x stands for the percent of trials in which the font was seen: $20x + .5(20 – 20x) = .6(20)$. "20" is the number of trials in each condition, and ".6" is average font accuracy. The equation specifies that the percent of trials in which the font is seen plus 50% of the remaining trails (viz., trials in which the font would be guessed accurately due to chance) equals the percent of trials in which the font is reportable, based on the observed font accuracy rate. For instance, if you know 10% of questions on a 20-question true/false test and guess on the rest, then the total number of questions you would be expected to answer correctly is .10*20, or 2, plus 50% of the remaining questions (9, since 20 – 2 = 18, and 19/2 = 9). This means you would be expected to answer 11 questions correctly, or 55% of the questions. Applied to Experiment 1, solving for the equation yields that participants saw the font of the words 20% of the time. This is the value that would yield 60% accuracy overall, taking into account the likelihood of correctly guessing the font by chance.

Assuming that font is visible on 20% of trials, an increase in font report accuracy would be expected in the Familiar Match condition if the familiar matching silhouettes allowed for retrieval of the words (since, as a conservative estimate, 20% of the words recovered as a result of the familiar matching silhouettes should have visible font).

We also looked at font report accuracy among participants who accurately reported the font on at least 60% of trials. We call these participants "high performers" (N = 20). High performers accurately reported the font of the words on average in 70% of trials. Since these participants were able to see the font on a higher percentage of trials (approximately 40%[2]) in all conditions, the expected percent of trials in which participants would be expected to accurately report on the font in the Familiar Match condition is even greater. However, even for High Performers, font report accuracy was no higher in the Familiar Match condition than any other condition.

We therefore think the experiment is some evidence that the episodic retrieval explanation is not the proper explanation of Sanguinetti & Peterson's unmasking effect, since there was no increase at all in font accuracy in the Familiar Match condition.

Interpreted within the unmasking framework, the experiment suggests that features of word (like their font) are not necessarily unmasked along with their semantics. On the hypothesis confirmation explanation of the unmasking phenomenon, the fact that the font is not unmasked with the semantics of the word is intriguing. A few possible explanations exist. Perhaps since the silhouettes are only related to the words semantically, only the semantics of the word is confirmed, and font information is lost. Thus, the unmasked words appear do not appear in any particular font.

Another intriguing but speculative explanation for this effect is that the hypothesis undergoing confirmation is not the token word presented prior to the random letter mask, but rather, a new semantic entity posited to explain the semantic activation from the word and the matching silhouette. It

2 If $20x + 5(20 - 20x) = .7(20)$, $x = .4$

is perhaps perplexing that the brain would hypothesize the persisting semantic information in the visual stream to originate from a single source, given the featural dissimilarity of the word and silhouette. Perhaps as a "best guess" as to the source of the persisting semantic information, the brain posits a semantic entity that has the form of a written word (consistent with the activation of word form from the word) but that lacks commitment to more fine-grained features of the word such as its font.

Before strong conclusions can be made, another possible explanation for the unmasking effect must be ruled out. Participants' difficulty in reporting the font of the word is consistent with the inference explanation of the increased word perception in the *Familiar Match* condition. According to the inference explanation, participants infer the identity of the word by noting the match between the Familiar silhouette and a possible interpretation of the incomplete "letter string" extracted from the fused word-random letter object. If the word has been visually fused with the random letter string, the font of the word may be impossible to identify because the random letter string was always in regular font, confusing the perception of the font of the word. In Experiment 2a, we aimed to further explore the unmasking effect and distinguish between the inference explanation and the hypothesis confirmation explanation offered by S&P for the results.

Experiment 2a

In Experiment 2a we tested a key prediction of the inference explanation: that silhouette recognition is sufficient for the unmasking effect. In Experiment 2, the silhouettes were presented for 70ms rather than 283ms. At 70ms, the silhouettes were expected to be equally identifiable. (This prediction was tested in Experiment 2b). On the inference explanation, then, unmasking should not be reduced, since conscious identification of the silhouettes should be sufficient for the effect to emerge.

Experiment 2a also allowed us to examine the effect of reduced silhouette exposure duration on the unmasking effect. The Experimental Novel silhouettes, containing familiar objects on the outer borders of the novel central objects, did not allow for unmasking in Sanguinetti and Peterson's (2013)

experiment nor in Experiment 1. This is perhaps because the outer shapes were not confirmed as objects in the display; only the central region was assigned object status. With reduced exposure duration, the familiar silhouettes may behave more like the experimental novel silhouettes because there would be less iterations confirming the presence of the familiar objects in the silhouettes.

Participants

56 undergraduate students from the University of Arizona participated in this study after giving informed consent. Participants completed the experiment to fulfill a requirement for their introductory psychology course. Participants all reported normal or corrected-to-normal visual acuity. Data from an additional 6 participants were excluded because they reported seeing the objects on the groundsides of the Experimental Novel silhouettes, leaving 50 participants in the main analyses. There were much less seers in this experiment than in Experiment 1, in which there were 22. This is likely due to the shorter silhouette exposure duration.

Design and Procedure

The experimental design was identical to the design of Experiment 1 except for the following modifications. All words were in regular font. The silhouettes were presented for 70 ms instead of 283 ms. Participants were asked to report only on whether they saw a word on each trial. Thus, the experimental design was identical to Sanguinetti & Peterson's original experiment except the silhouette exposure duration was decreased from 175ms to 70ms, and participants were only asked to report whether they saw a word on each trial; they were not asked to categorize the silhouettes or to indicate whether the word matched the silhouette on each trial.

Results

Figure 8 shows the mean and standard error bars for the percentage of trials in which participants reported seeing a word for each of the five within-subject conditions. A one-way ANOVA with the factor of condition (*Familiar Match*, *Familiar Mismatch*, *Experimental Match*, *Experimental*

Mismatch, and *Control*) revealed a significant effect of condition, $F(4,49) = 8.99$, $p < .001$. A post-hoc Least Significant Difference (LSD) test corrected for multiple comparisons revealed that responses of seeing a word were significantly higher in the *Familiar Match* condition (M=.28, SE=.04) than all other conditions. Thus, the unmasking effect was replicated.

However, the unmasking effect was reduced compared to Sanguinetti and Peterson (2013). Mean word perception was 39% in Sanguinetti & Peterson's original experiment in the *Familiar Match* condition, whereas it was only 28% in the present experiment. This indicates that the unmasking effect, although present with a 70ms silhouette exposure duration, is reduced.

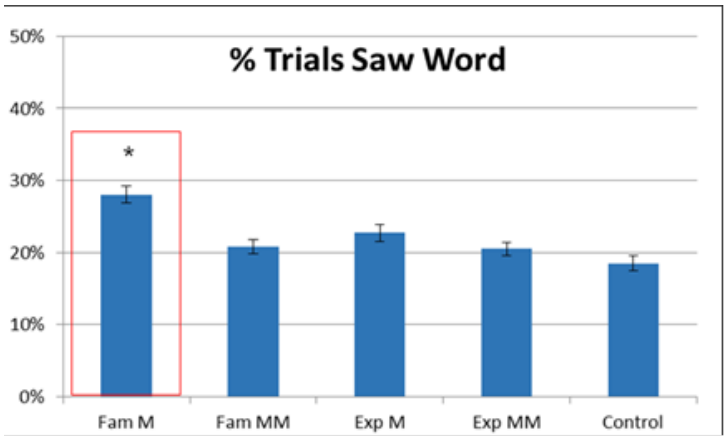

Figure 8 Word perception, Experiment 2a. The unmasking effect was reduced from 39% word perception in the Familiar Match condition in Sanguinetti & Peterson (2013) to 28% word perception in the present study, which used a 70ms silhouette exposure duration.

Discussion

Experiment 2a indicates that the degree of unmasking depends on silhouette exposure duration. With a shorter exposure duration, the unmasking effect was reduced in the *Familiar Match* condition. Experiment 2a also speaks to the relationship between silhouette identification and unmasking. We expect that at the 70ms exposure, participants were equally able to identify the silhouettes as they were at the 175ms exposure duration used by Sanguinetti & Peterson. Even so, the unmasking effect was reduced. This is counter to the prediction of the inference explanation, on which there is no reason to suspect that reducing silhouette exposure would reduce the unmasking effect, as long as silhouette identification is unaffected.

The reduced unmasking effect is consistent with the inability of the Experimental Novel silhouettes to facilitate word perception. The objects on the groundsides of the Experimental Novel silhouettes are processed semantically, even when unseen. The fact that they do not facilitate word perception indicates that semantic activation consistent with the word is not sufficient for unmasking—the familiar region must be sufficiently confirmed as an object in the display for unmasking to occur. The reduced unmasking associated with the shorter silhouette exposure duration supports the claim that sufficient processing confirming the presence of familiar regions as objects in the display is required for unmasking.

For a final experiment, we tested the assumption that silhouette recognition was equally accurate at a 70ms exposure duration as a 175ms exposure duration. To test this, participants categorized the silhouettes as Novel or Familiar as they did in Sanguinetti and Peterson(2013). If categorization is equally accurate, it would suggest that silhouette identification is equally good at a 70ms exposure as it is at 175ms.

Experiment 2b

Participants

25 undergraduate students from the University of Arizona participated in this study after giving informed consent. Participants completed the experiment to fulfill a requirement for their introductory psychology course. Participants all reported normal or corrected-to-normal visual acuity. In this experiment, no participants reported seeing the objects on the groundsides of the Experimental Novel silhouettes. The data of all participants was thus included.

Design and Procedure

The experimental design was identical to the design of Experiment 2a except rather than reporting on whether they saw a word on each trial, participants categorized the silhouettes as novel or familiar. The exposure duration of the silhouettes was 70ms. Participants were not asked to respond

about the words, although they were told, as they were in the previous experiments, that most of the trials would contain words.

Results

Participants were accurate in categorizing the silhouettes as familiar or novel. Accuracy did not drop below 74% (SE = 2%) in any condition. *Figure 9* compares categorization accuracy in the present experiment, which used a 70ms silhouette exposure duration, to accuracy in Sanguinetti and Peterson(2013), which used a 175ms silhouette exposure duration. A 2-way ANOVA with the factors of condition *(Familiar Match, Familiar Mismatch, Experimental Match, Experimental Mismatch, and Control*) and the between-subject factor of Exposure Duration (175ms vs. 70ms) revealed no significant effects of either factor and no interaction. This indicates that silhouette categorization

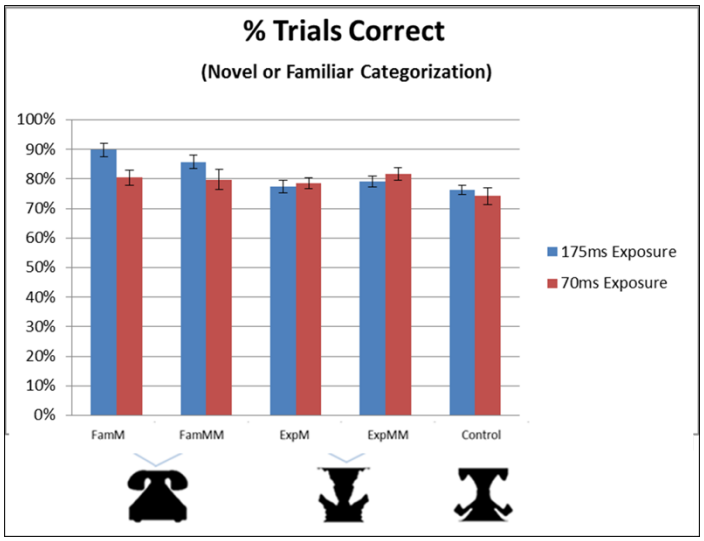

Figure 9. Comparison of silhouette categorization accuracy in Sanguinetti & Peterson (2013), represented by blue bars, and Experiment 2b, represented by red bars. Accuracy was high in both experiments and did not differ between them.

accuracy was no higher in Sanguinetti & Peterson's experiment than in the present experiment.

Discussion

Experiment 2b provides evidence that silhouette identification is not sufficient for the degree of unmasking observed in Sanguinetti & Peterson's original experiment. Silhouette identification was

equally good at the 70ms exposure duration as at the 175ms exposure duration, but unmasking in the Familiar Match condition was reduced by 11% when the silhouettes were displayed for only 70ms.

A limitation of the experiment is that the silhouettes only had to be categorized as Novel or Familiar. Subjects did not need to identify the basic-level category of the silhouettes (e.g., dog). Perhaps subjects did not in fact identify the basic-level category of the silhouettes in order to categorize them as Novel or Familiar. This would render the present experiment an ineffective test of whether basic-level silhouette identification is sufficient for the unmasking effect. We accept this worry as valid; subjects may have categorized the silhouettes based on a sense of the silhouette being familiar or novel without actually identifying the particular objects in the silhouette. However, one might accept some diminishment in accuracy if participants employed a general familiarity or novelty signal in categorizing the silhouettes, yet no decrease in accuracy was observed.

Because the omnibus ANOVA was not significant, comparisons of accuracy for each exposure duration (175ms vs 70ms) within conditions was not performed. However, it should be noted that accuracy was approximately 10% lower in the *Familiar Match* condition at the 70ms exposure duration than the 175ms exposure duration. Because this is the condition for which unmasking was reduced at 70ms, this finding warrants consideration. However, we think that if there was indeed a reduction in accuracy in silhouette categorization in the *Familiar Match* condition at the 70ms exposure duration, this was not in virtue of the duration alone, but rather, because the increased word perception that occurred at the 175ms exposure aided in silhouette categorization. For instance, the experience of a word matching the silhouette may facilitate categorization of the silhouette as familiar. We think the inference that exposure duration alone does not explain the effect is warranted because silhouette categorization accuracy in the other conditions was not affected by the decreased exposure duration. Thus, it is highly likely that the increased word perception at the 175ms exposure explains the higher categorization accuracy at the 175ms exposure in the Familiar Match condition.

General Discussion

We presented three experiments aimed at discerning the explanation for a semantic unmasking phenomenon. The results are consistent with Sanguinetti & Peterson's (2013) original interpretation of the unmasking effect involving feedback processing and hypothesis confirmation of perceptual stimuli, and do not support the episodic retrieval or inference interpretations. In Experiment 1, although participants saw more words in the Familiar Match condition than in the other condition, there is no evidence that they saw the font of the words that were unmasked. If the Familiar silhouettes simply caused participants to recall an earlier episode of seeing the words, then they would be expected to be able to report on the font of the words. Experiments 2a and 2b demonstrated that conscious identification of the semantically related silhouettes is not sufficient for the unmasking effect. It is therefore unlikely that participants inferred the words they had seen from the objects depicted in the silhouettes.

As discussed, a weakness of Experiment 1 is that font report accuracy was generally low, even with an exposure duration of 83ms (pre- and post-masked). Given that font reports were not much better than expected by chance, the increase in font reports expected to occur in the familiar match condition (given that the words were unmasked or retrieved along with their font) is likely to be small. Although we did not observe any difference in font report accuracy at all between the Familiar Match condition and other conditions in Experiment 1, the fact that the expected difference is small warrants caution in interpreting this result. It is possible that a difference would emerge with a larger sample size. Participants who accurately reported the font on at least 60% of trials (i.e., “high performers) also were not more accurate at reporting the font in the Familiar Match condition than in the other conditions. However, even for high performers the expected increase in font accuracy in the Familiar Match condition is likely fairly small.

Given these limitations of these analyses, it is worth conducting additional experiments to improve confidence in the conclusions of Experiment 1. One possible modification is to increase the exposure duration of the words. The words were presented for a longer than intended due to a programming mistake, which worked to our advantage, since the font of the words was more difficult to report than expected. However, it is worth increasing the duration more to better ensure that the font is visible along with the word when the word is reportable. This would likely increase word perception even more, but would have the advantage of increasing font report accuracy in all conditions. If font report accuracy is fairly high (i.e., as high as it was for "high performers" for all participants, or higher), then if font report accuracy is higher in the Familiar Match condition than the other conditions, this could not be attributable to the font of the word failing to be perceived along with its semantics. It would therefore be a more effective test of the episodic retrieval hypothesis.

It also might be advantageous to repeat Experiment 1 using a feature that is easier to perceive than font, such as color, to ensure that the feature is perceived along with the semantics of the word when participants report seeing a word. However, there are pitfalls that must be watched out for with the use of any given alternative feature. For instance, the feature must not be so obvious that is reportable on every trial, even when words are not seen. Color might fall prey to this worry.

With respect to the paradigm used in Experiment 1, proponents of the episodic memory interpretation might also make the following general theoretical objection. Nothing can be inferred from there being no increase in font report accuracy (or accuracy at reporting any other feature) in the Familiar Match condition for the following reason. Words that are forgotten would be less likely to be encoded along with their font (or other features) than words that were sufficiently perceived upon presentation to be made available to working memory. Because the extra words seen in the Familiar Match condition would be classified as "forgotten" (because they were not recalled until the appearance of matching silhouettes), those words in particular would not be expected to have been

encoded along with their font (or other features), and therefore, their font would not be recalled along with their semantics. However, accepting this proposal entails acknowledging that the episodic retrieval hypothesis is in principle untestable using the paradigm under consideration (i.e., asking participants to report on a feature of the word). One could always maintain that with respect to the "extra" words perceived in the Familiar Match condition, if their features are not reportable along with their semantics, this is simply because the features had not been encoded with the semantics for those particular words. Given the paradigm under consideration seems to be a *prima facie* fair test of the episodic retrieval hypothesis, this post-hoc reason for discounting the results seems dubious.

Experiment 1 provides some evidence that unmasking does not necessarily preserve all of the information contained in the original representation. In other words, it suggests that some components of stimuli (such as their semantics) might be unmasked while others (such as their font) are not.

Experiments 2a and 2b showed the conscious recognition of the silhouettes is insufficient for the unmasking effect, and furthermore, that unmasking defends on the extent of silhouette processing. Further research might address exactly how much silhouette processing must occur for unmasking.

The present series of experiments was designed to test a few alternative explanations for Sanguinetti & Peterson's (2013) unmasking phenomenon: the episodic retrieval explanation and the inference explanation. These are not the only possible alternative explanations for Sanguinetti and Peterson's results. That is, although the results are consistent with Sanguinetti and Peterson's interpretation involving feedback processing, they do not necessarily show that this is the correct explanation for the results.

Another possible explanation was suggested by Steven Palmer when this work was presented at the Vision Science Society conference in 2015. Palmer suggested that the unmasking effect occurs because the semantically related silhouettes cause activation at the lexical node for the previously

presented word. The activation from the silhouette at the node combines with the activation from the previously presented related word, rendering the activation sufficient for awareness of the word.

This explanation is a viable alternative. We think that the following considerations count against this alternative, however (or at least, renders it incomplete). It is not at all obvious or clear that accumulation of semantic information at a node (including strong activation for the orthographic representation at the node) in itself contributes to a stimulus being phenomenally conscious. In semantic priming and repetition priming, the effect of semantic information on lexical decision times might be the result of the facilitation of word *recognition* (or other kinds of processing of the word stimulus), rather than an effect on the strength of the word *percept* (i.e., the extent to which the word is phenomenally conscious). Lexical decision tasks do not distinguish between these alternatives. The present experiments differ in that participants do not merely make a response or speeded decision about the word, but are asked to report specifically whether the word is phenomenally conscious. One worry is that participants sometimes reported seeing a word even if they simply guessed that word was present, or if it seemed to them that a word was present without having actually perceived this. To better ensure that the words were phenomenally conscious, it would be advantageous to include a rating scale for word visibility (e.g., “not visible at all” to “very visible”).

Furthermore, even if semantic context indeed contributes to the phenomenal consciousness of words (and other semantically relevant stimuli), then this itself might be best explained by the hypothesis confirmation/iterative processing model. Such findings described loosely in terms of activation at nodes might be reinterpreted as involving hypothesis confirmation about the stimuli in question, with the extent of semantic information mattering for the likelihood of any given hypothesis about the presence of the semantic item to be confirmed.

Conclusion

In Experiment 1, we showed that an attended feature of the words (i.e., their font) is not unmasked along with the semantics of the words—unmasked words appear in no identifiable font. This is evidence against the episodic retrieval hypothesis, since on the episodic retrieval hypothesis the words would be expected to be retrieved along with their font as the font would be part of the original episode of seeing the words. In Experiments 2a and 2b, we show that conscious identification of the objects in the silhouettes is insufficient for the unmasking effect—reduction in silhouette exposure duration reduces the unmasking effect (Experiment 2a) even though silhouettes are equally recognizable (Experiment 2b). This rules out the inference explanation for the unmasking effect, since on this explanation unmasking would depend on whether the silhouettes are recognized and can therefore be employed in conscious inference. Sanguinetti & Peterson's original interpretation of the unmasking effect holds up against the tests presented in this thesis.

References

Bouvier, S., & Treisman, A. (2010). Visual feature binding requires reentry. *Psychological Science, 21*(2), 200-204.

Breitmeyer, B. G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing. *Psychological Review, 83*(1), 1.

Breitmeyer, B. G., Rudd, M., & Dunn, K. (1981). Metacontrast investigations of sustained–transient channel inhibitory interactions. *Journal of Experimental Psychology: Human Perception and Performance, 7*(4), 770.

Briscoe, G., Dember, W. N., & Warm, J. S. (1983). Target recovery in visual background masking: No clear explanation in sight. *Journal of Experimental Psychology: Human Perception and Performance, 9*(6), 898.

Chen, H., & Wyble, B. (2015). Amnesia for object attributes: Failure to report attended information that had just reached conscious awareness. *Psychological Science, 26*(2), 203-210.

Corthout, E., Uttl, B., Walsh, V., Hallett, M., & Cowey, A. (1999). Timing of activity in early visual cortex as revealed by transcranial magnetic stimulation. *Neuroreport, 10*(12), 2631-2634.

Dell'Acqua, R., Pascali, A., Jolicoeur, P., & Sessa, P. (2003). Four-dot masking produces the attentional blink. *Vision Research, 43*(18), 1907-1913.

Dell'Acqua, R., Pascali, A., Jolicoeur, P., & Sessa, P. (2003). Four-dot masking produces the attentional blink. *Vision Research, 43*(18), 1907-1913.

Di Lollo, V. (1980). Temporal integration in visual memory. *Journal of Experimental Psychology: General, 109*(1), 75.

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General, 129*(4), 481.

Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science, 8*(2), 135-139.

Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences, 4*(9), 345-352.

Francis, G., & Hermens, F. (2002). Comment on" competition for consciousness among visual events: The psychophysics of reentrant visual processes"(di lollo, enns & rensink, 2000).

Kahneman, D. (1968). Method, findings, and theory in studies of visual masking. *Psychological Bulletin, 70*(6p1), 404.

Keysers, C., & Perrett, D. I. (2002). Visual masking and RSVP reveal neural competition. *Trends in Cognitive Sciences, 6*(3), 120-125.

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. *Philosophical Transactions of the Royal Society of London.Series B, Biological Sciences, 362*(1481), 857-875.

Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). A new neural framework for visuospatial processing. *Nature Reviews Neuroscience, 12*(4), 217-230.

Kulikowski, J., & Tolhurst, D. (1973). Psychophysical evidence for sustained and transient detectors in human vision. *The Journal of Physiology, 232*(1), 149-162.

Kulikowski, J., & Tolhurst, D. (1973). Psychophysical evidence for sustained and transient detectors in human vision. *The Journal of Physiology, 232*(1), 149-162.

Lamme, V. A., Super, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology, 8*(4), 529-535.

Lamme, V., Zipser, K., & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *Cognitive Neuroscience, Journal of, 14*(7), 1044-1053.

Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 15*(2), 1605-1615.

Lleras, A., & Moore, C. M. (2003). When the target becomes the mask: Using apparent motion to isolate the object-level component of object substitution masking. *Journal of Experimental Psychology: Human Perception and Performance, 29*(1), 106.

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology, 22*(1), 1-35.

Marcel, A. J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15*(2), 197-237.

Neill, W. T. (1997). Episodic retrieval in negative priming and repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*(6), 1291.

Ögmen, H., Breitmeyer, B. G., Todd, S., & Mardon, L. (2006). Target recovery in metacontrast: The

effect of contrast. *Vision Research, 46*(28), 4726-4734.

Peterson, M. A. (1994). Object recognition processes can and do operate before figure-ground organization. *Current Directions in Psychological Science,* , 105-111.

Peterson, M. A., & Gibson, B. S. (1993). Shape recognition inputs to figure-ground organization in three-dimensional displays. *Cognitive Psychology, 25*(3), 383-429.

Peterson, M. A., & Gibson, B. S. (1994). Must figure-ground organization precede object recognition? an assumption in peril. *Psychological Science, 5*(5), 253-259.

Peterson, M. A., & Gibson, B. S. (1994). Must figure-ground organization precede object recognition? an assumption in peril. *Psychological Science, 5*(5), 253-259.

Peterson, M. A., & Kim, J. H. (2001). On what is bound in figures and grounds. *Visual Cognition, 8*(3-5), 329-348.

Peterson, M. A., & Skow, E. (2008). Inhibitory competition between shape properties in figure-ground perception. *Journal of Experimental Psychology: Human Perception and Performance, 34*(2), 251.

Reiss, J. E., & Hoffman, J. E. (2007). Disruption of early face recognition processes by object substitution masking. *Visual Cognition, 15*(7), 789-798.

Sanguinetti, J. L., & Peterson, M. A. (2013). Unmasking the mask: Semantic similarity produces disinhibition in a masked priming paradigm. *Journal of Vision, 13*(9), 47-47.

Shapiro, K. L., Raymond, J., & Arnell, K. (1997). The attentional blink. *Trends in Cognitive Sciences, 1*(8), 291-296.

Shapiro, K. L., Raymond, J., & Arnell, K. (1997). The attentional blink. *Trends in Cognitive Sciences, 1*(8), 291-296.

Spencer, T. J., & Shuntich, R. (1970). Evidence for an interruption theory of backward masking. *Journal of Experimental Psychology, 85*(2), 198.

Tenkink, E. (1983). Recovery at short intervals between masking flashes. *Vision Research, 23*(12), 1693-1698.

Trujillo, L. T., Allen, J. J., Schnyer, D. M., & Peterson, M. A. (2010). Neurophysiological evidence for the influence of past experience on figure-ground perception. *Journal of Vision, 10*(2), 5.1-21.

Turvey, M. T. (1973). On peripheral and central processes in vision: Inferences from an information-processing analysis of masking with patterned stimuli. *Psychological Review, 80*(1), 1.

Weisstein, N., Ozog, G., & Szoc, R. (1975). A comparison and elaboration of two models of metacontrast. *Psychological Review, 82*(5), 325.