Journal of Experimental Psychology: General
2008, Vol. 137, No. 2, 324–347

Copyright 2008 by the American Psychological Association
0096-3445/08/$12.00 DOI: 10.1037/0096-3445.137.2.324

Immediate Priming and Cognitive Aftereffects

David E. Huber
University of California, San Diego

Three forced-choice perceptual word identification experiments tested the claim that transitions from positive to negative priming as a function of increasing prime duration are due to cognitive aftereffects. These aftereffects are similar in nature to perceptual aftereffects that produce a negative image due to overexposure and habituation to a stimulus. Each experiment tested critical predictions that come from including habituation in a dynamic neural network with multiple levels of processing. The success of this account in explaining the dynamics of repetition priming, associative-semantic priming, and forward masking effects suggests that habituation is a useful mechanism for reducing source confusion between successively presented stimuli. Implications are considered for visible persistence, repetition blindness, attention-based negative priming, attentional blink, inhibition of return, the negative compatibility effect, affect priming, and flanker preview effects.

Keywords: repetition priming, semantic priming, masking, neural networks

More than 30 years ago, Meyer and Schvaneveldt (1971) demonstrated that lexical decisions to associated word pairs (e.g., accurately stating that both *doctor* and *nurse* are valid words) are faster than lexical decisions to unrelated word pairs. Since then, the typical immediate priming paradigm has changed slightly, presenting first a prime word followed by a target word that either is briefly flashed and masked (perceptual identification) or remains onscreen while participants rapidly give a lexical decision or naming response. This paradigm is termed *immediate* or *short-term priming* because the advantage inferred by a related prime is relatively short-lived, occurring only when primes are presented on the same trial as targets (Joordens & Becker, 1997; Ratcliff, Hockley, & McKoon, 1985). The term *priming* is used because the first word "primes the pump" for a related target, yielding faster or more accurate responding.

Since its inception, immediate priming has been used as a measurement tool, with priming facilitations indicating the nature of perceptual similarity (Leopold, O'Toole, Vetter, & Blanz, 2001; Long, Toppino, & Mondin, 1992), semantic structure (McNamara, 1992, 1994; Ratcliff & McKoon, 1994), orthographic or phonemic similarity (Lukatela, Eaton, Lee, Carello, & Turvey, 2002; Lukatela, Frost, & Turvey, 1998; Meyer, Schvaneveldt, & Ruddy, 1974), response competition (Tipper & Driver, 1988; Yee, Santoro, Grey, & Woog, 2000), syntactic processing (Ferreira, 2003), affective evaluation (Murphy & Zajonc, 1993; Stapel & Winkielman, 1998), or implicit attitudes (Dasgupta, Greenwald, & Banaji, 2003; Greenwald, McGhee, & Schwartz, 1998), to name just a few examples. In this research, I take a different approach by studying the perceptual, memory, and decision mechanisms that support the basic immediate priming phenomena in all of these domains of study. From this perspective, immediate priming indicates basic aspects of cognition that may apply to a wide variety of tasks and judgments. By analogy, the difference between these approaches is similar to the difference between priming the pump to measure the properties of water (e.g., the lexicon) versus studying the pump itself to figure out how it works (e.g., cognitive processes).

Despite the ubiquitous use of immediate priming, the results are often varied, with some experiments finding facilitation in one situation and others failing to find facilitation or even finding reliable priming deficits in the same situation (Hochhaus & Johnston, 1996; Leopold et al., 2001; Long et al., 1992; Tipper & Driver, 1988). Providing one explanation for these inconsistencies, Huber, Shiffrin, Lyle, and Ruys (2001) examined perceptual identification with forced-choice testing, which allows separation of decisional and perceptual aspects of priming (see Figure 1 for an example of this paradigm as applied to the current experiments). The basic task of the observer was to attempt to identify the briefly flashed target word that appeared in the center of the screen. Rather than measuring naming accuracy for the target word amongst all possible words, performance was measured by means of a forced choice between the correct target word and an incorrect foil word. Target durations were placed at the perceptual threshold, and choice accuracy provided the measure of priming. Unlike most priming experiments, this paradigm included just as many trials that primed the incorrect foil word as trials that primed the correct target word. Participants were explicitly told that choosing on the basis of the prime was just as likely to produce the wrong answer as the correct answer in order to discourage strategies. Furthermore, these instructions were validated through trial-by-trial accuracy feedback.

With this paradigm, Huber et al. (2001) found clear evidence that the magnitude of priming and even the direction of priming can change as a function of how the primes are processed. In one set of conditions, termed *passive priming*, participants simply viewed briefly presented primes, and in another set of conditions, termed *active priming*, participants were given a task to complete in relation to the primes prior to the briefly flashed target. For the

This research was supported by National Institute of Mental Health Grants MH63993-01 and MH063993-04, National Science Foundation Grant KDI/LIS IBN-9873492, and Office of Naval Research Grant N00014-00-1-0246.

Correspondence concerning this article should be addressed to David E. Huber, Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109. E-mail: dhuber@psy.ucsd.edu

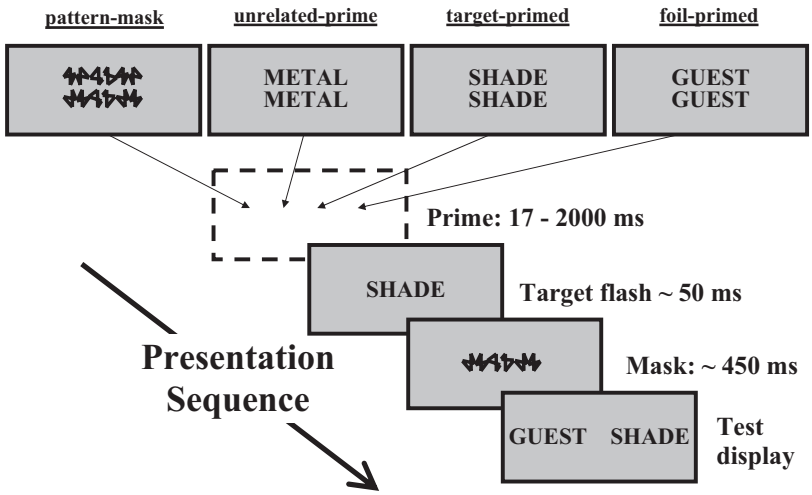


Figure 1. Sequence of events for the priming conditions in Experiment 1. Participants attempted to identify the briefly flashed target word, which was the only word that appeared in the center of the screen. Primes were presented for 17, 50, 150, 400, or 2,000 ms. Target flash durations were determined separately for each participant to place performance at 75% correct (perceptual threshold) as measured by accuracy in a forced choice (test display) between the target word and a foil word. Mask durations were set such that the total time between target flash onset and test display onset was 500 ms. Trial-by-trial feedback was provided to minimize strategic responding.

case of repetition priming of the target (e.g., the *target-primed* condition in Figure 1), performance was better when the target word was passively primed but worse with active priming. Furthermore, for the case of repetition priming the incorrect foil word (e.g., the *foil-primed* condition in Figure 1), the opposite pattern was observed, with deficits for passive priming switching to benefits with active priming. Subsequent experiments demonstrated that passive viewing of primes of different durations was sufficient to produce this change in the direction of priming, although other experiments demonstrated that responding to primes bolsters the effect (Huber, Shiffrin, Quach, & Lyle, 2002; Weidemann, Huber, & Shiffrin, 2005). The observation that the direction of priming is changeable places a cautionary note on any immediate priming experiment, suggesting that the visual, temporal, and attentional characteristics of prime presentations need to be carefully considered.

Perceptual Aftereffects, Cognitive Aftereffects, and the Type–Token Distinction

Another area of research producing effects that vary in magnitude and even direction as a function of exposure is the study of low-level visual perception. Brief or minimal displays result in a positive image, but there are many examples of perceptual aftereffects that eliminate this positive image or even produce negative images following extended viewing. Examples include but are not limited to the phenomena of color opponency (Shively, 1973), Troxler fading (Kotulak & Schor, 1986), or the "Jesus Illusion" shown in Figure 2. For these phenomena, the longer that a stimulus is viewed (e.g., maintained fixation on Figure 2), the more strongly the opposite pattern will subsequently emerge in comparison to some new stimulus (e.g., switching one's view to a white wall).

Aftereffects such as the Jesus Illusion are clearly perceptual in nature. However, the claim in my work is that similar aftereffects exist not only for low-level perception but also for higher forms of perception and identification. Because there is substantial disagreement regarding where to place the dividing line between perception and higher forms of identification (e.g., language and semantics), the phrase *cognitive aftereffect* is used in referring to this larger class of aftereffects that extends beyond low-level perception. This is not to say that cognitive aftereffects involve cognition per se or are due to a strategic thought process. Instead,



Figure 2. The "Jesus Illusion" (origin unknown) demonstrating perceptual aftereffects. After maintaining fixation on the lowest of the four vertical dots for approximately 30 s, stare at a white wall while blinking to see the afterimage of Jesus.

the claim is simply that the dynamics commonly observed for low-level perception can exist for other types of processing. In this manner, *cognitive aftereffect* is used as a catch-all phrase for all sorts of aftereffects that exist in the process of identification. In my work on immediate priming, I assume that the mechanisms supporting perceptual aftereffects are functionally similar to the mechanisms supporting higher level cognitive aftereffects and priming. Furthermore, aftereffects are a feature rather than a bug—habituation produces aftereffects to reduce source confusion in processing the rapidly changing stream of events.

Source confusion is the tendency to mix up what happened when and where. To understand the nature of source confusion, consider a detector that responds to a certain feature in the world regardless of when or where that feature appears. For instance, through practice as a child, you learned to tell the difference between a *P* and an *R* presumably because your *R* detector is highly sensitive to the presence of a diagonal line protruding from a rounded top. Furthermore, your ability to detect *R*s works well regardless of when or where you see an *R*, how large it is, the font color, and so on. Presumably you did not need to learn every conceivable version of *R* separately. Instead, you have developed a more centralized detection process that readily generalizes to many situations. Such a generalized detector is termed a *type* detector because it is sensitive to the generic identification of a stimulus (e.g., any *R*) rather than specific instantiations of a stimulus (e.g., an *R* that is in 22-point gray Times Roman font and appears 0.75 degrees to the left of the fovea at 3:21 p.m.). However, there is a downside to this generalized type detection ability. Because type detection responds to any *R*, it suffers from source confusion and has trouble discriminating between an *R* that was flashed in the word *PRIME* 100 ms ago versus one that was flashed 50 ms ago in the word *TARGET*. These separate specific instantiations of *R*s are referred to as *tokens*. My claim is that negative aftereffects provide a natural and automatic mechanism that serves to reduce source confusion between separate tokens that occur in rapid succession.

One solution to the problem of source confusion is to build many token detectors to identify every possible instantiation. However, this solution is not parsimonious and results in a combinatorial explosion of type–token pairs, with several tokens needed for each type. Furthermore, this solution poses a problem for perceptual learning and requires modifying multiple token detectors in concert. Yet in many situations, such as deciding how many occurrences of an *R* were seen, it is necessary to individuate between separate tokens. This type–token distinction lies at the heart of the phenomenon termed *repetition blindness*, in which a second occurrence of a word in a rapidly presented sentence is often missed (Hochhaus & Johnston, 1996; Kanwisher, 1987; Kanwisher, Kim, & Wickens, 1996). My explanation of repetition blindness, as well as many other experiments that involve immediate effects of one stimulus on a subsequent stimulus, is that people's perceptual systems have finely tuned temporal dynamics in the type system that minimize source confusion by reducing responses after sufficient identification. In other words, once you have seen one *R* token, the *R* detector enters into a transient refractory period to minimize blending with subsequent presentations. However, this solution comes with a cost and produces certain misperceptions. In particular, repetition blindness and other aftereffects are signature misperceptions of this tokenization through temporary suppression.

According to this theory, aftereffects are the attempt to parse the stream of perception into separate events through the discounting of previously identified objects, with this discounting producing negative priming.[1] Therefore, aftereffects produce a perceptual system that is primarily sensitive to the onset of new events. Next, I review several theoretical models that describe negative priming as observed in perceptual identification tasks. To preview the discussion, note that all three of the extant models appeal to the discounting of primed items, but only the neural habituation account of Huber and O'Reilly (2003) provides explicit predictions regarding the timing of events. Previous work on priming dichotomously examined long versus short primes or subliminal versus supraliminal primes, but in the current experiments, I parametrically varied prime duration to test specific predictions of this dynamic model.

Immediate Priming: Models That Produce Negative Priming Through Discounting

ROUSE: Bayesian Discounting of Features

The *responding optimally with unknown sources of evidence* (ROUSE) model accounts for a change from positive to negative priming through the offsetting mechanisms of source confusion and discounting. With source confusion (unknown sources of evidence), the perceptual system is mixed up regarding what happened when and where, and features of the primes are mistaken for the target (e.g., a *T* was seen, but it was unclear which perceptual event it belonged to). With discounting (responding optimally), features that are known to have been primed are assigned a lower level of evidence (i.e., weighted less heavily) because the prime rather than the target may have been the source. For instance, if it a *T* is perceived and it is known that the prime contained a *T*, then this still counts in favor of a choice alternative such as *TREAT* but less so than if the prime did not contain a *T*. This calculation is achieved using a rational analysis of the task (Schooler & Anderson, 1997), resulting in the same mathematical phenomena referred to as "explaining away" in the study of Bayesian belief networks (Castillo, Gutierrez, & Hadi, 1997). Therefore, discounting is the attempt to normatively (i.e., optimally) correct for source confusion. The switch between positive and negative priming corresponds to too little versus too much discounting as compared with the actual level of source confusion. The model is highly successful at accounting for situations that give rise to positive or negative priming as a function of prime–target similarity, similarity between the choice words, or target duration (Huber, Shiffrin, Lyle, & Quach, 2002). Nevertheless, the ROUSE model does not specify the relationship between discounting and prime duration.

[1] There may be some confusion between the term *negative priming* as it is used here in referring to lower accuracy in perceptual identification versus the specific paradigm of negative priming, which involves slower responding to previously ignored stimuli (Tipper, 1985). In the current context, the term *negative priming* simply means that the difference in performance between priming the correct answer versus priming the incorrect answer is negative in value.

Multinomial Model: Discounting Entire Words Rather Than Features

Demonstrating the robustness of discounting as an explanatory concept, Ratcliff and McKoon (2001) implemented source confusion and discounting in an all-or-none manner for entire words rather than features. This was accomplished with a simple multinomial model. With all-or-none discounting (e.g., with some probability, the observer decides not to choose the primed word on a particular trial), this leaves open the possibility that these effects are strategic. Furthermore, the multinomial model is somewhat descriptive because it fails to include a mechanism to handle changes as a function of similarity between words. In contrast, the ROUSE model manipulates similarity through the proportion of shared features, and this produced critical predictions regarding situations in which negative priming switches back to positive priming. Besides a lack of predictive power in regard to similarity, the multinomial model likewise fails to specify the relationship between prime duration and the direction of priming.

Neural Model: Discounting Through Habituation

Huber and O'Reilly (2003) developed a dynamic extension of the ROUSE model that uses neural habituation to explain the switch from positive to negative priming with prime duration. Because the reported experiments test key predictions of this dynamic model, the workings of this model are next considered in some detail. At its core, this model uses the mechanism of short-term habituation, which is ubiquitous in neural processing. For instance, single cell recordings consistently reveal spike rates (i.e., the number of action potentials per second) that quickly reach a maximum but then fall off to a low asymptotic rate (Duysens, Orban, Cremieux, & Maes, 1985). Neural habituation is studied by driving sending cells at different rates while recording from receiving cells. Experiments of this sort were the origin of the term *synaptic depression* (Abbott, Varela, Sen, & Nelson, 1997; Tsodyks & Markram, 1997) for the observation that the effect of the sending cell as measured at the receiving cell diminishes over time despite a constant firing rate. With synaptic depression, resources critical to effective signaling between cells are depleted as a function of recent activity in those cells. Many potential mechanisms are responsible for synaptic depression, including depletion of the vesicles necessary for neurotransmitter release, depletion of the calcium needed for vesicle release, failure of neurotransmitter reuptake, postsynaptic desensitization, and the buildup of slow calcium currents that directly inhibit the receiving cell. Abbott et al. formalized habituation through synaptic depression in a mathematical expression with one parameter for the rate of depletion and a second for the rate of recovery. Huber and O'Reilly used their equations to derive a version of neural habituation that is appropriate to levels of activation rather than the more complicated specifics of spiking cells.

Similar to real neurons, the model that Huber and O'Reilly (2003) developed contained a threshold for activation such that simulated cells do not produce any output until the threshold is crossed. This has implications for neural habituation because it is the above-threshold activity that drives depletion. Nevertheless, activation and depletion are dynamic processes that are always in operation, with the existence of any above-threshold activation serving to constantly drive depletion. In terms of mathematical instantiations of these processes, one dynamic update equation (i.e., differential equation) dictates the manner in which presynaptic activity changes at every millisecond, and a second update equation dictates the manner in which depletion changes at every millisecond. The multiplication of presynaptic activity and the level of available nondepleted synaptic resources yields the dynamically varying output activation that is sent to a receiving cell.

It may appear that activation and depletion are simply two sides of the same coin (as would be the case if a particular level of activation determined a set level of depletion), but in truth they are somewhat dissociable. For instance, these two processes do not operate on the same time scale; both behavioral data and cellular data suggest that recovery from depletion is relatively slow (e.g., hundreds or thousands of milliseconds rather than tens of milliseconds). Therefore, first a peak output activation is achieved due to the rapid buildup of presynaptic activity, but next, as more and more habituation accrues that is slow to recover, output activation falls to a low but above-baseline level. This being the case, depletion lasts longer than the effect of lingering activation from a prime (i.e., on the order of seconds rather than hundreds of milliseconds). Because depletion can outlive activation, depletion produces perceptual and cognitive aftereffects. By analogy, consider the process of eating versus the process of digestion, which likewise operate on different time scales. In this analogy, a sense of fullness is the product of eating and digestion, just as output activation is the product of presynaptic activity and the level of depletion. One can eat quite quickly and feel full in just a few minutes, and yet it may take many hours for digestion to remove that full feeling even though digestion is always in operation. As a result, eating behavior tends to oscillate, resulting in discrete meals rather than the slow rate of eating that would be required to keep fullness at a constant level. Similarly, presentation of a perceptual stimulus results in rapid activation but then habituation through depletion, resulting in a nonmonotonic rise then fall in the output activation that is sent to other cells.

Synaptic depression is the dominant form of short-term plasticity and is one of the largest responses in neural processing, often representing a magnitude order of change over the course of a few hundred milliseconds (S. B. Nelson, Varela, Sen, & Abbott, 1997). Huber and O'Reilly (2003) proposed that synaptic depression serves a critical function in perceptual processing, producing dynamic discounting in an attempt to overcome source confusion. Because higher level perception uses neurons with large receptive fields[2] (e.g., Tanaka, 1996), this results in source confusion when different stimuli are presented one after the other. Thus, synaptic depression may serve to parse the stream of perceptual events, thereby reducing interference from prior events in the identification of current events.

Figure 3 presents the three-level neural network proposed by Huber and O'Reilly (2003) to handle word perception and immediate word-priming phenomena. Except for the addition of synaptic depression (habituation), the dynamics of the simulated neurons

[2] The term *receptive field* refers to the specific patterns of visual input to which a cell responds most strongly, and the *size* of the receptive field refers to the visual area within which the preferred pattern produces a response. For instance, consider a cell that responds to the presentation of an *R*. If that cell has a large receptive field, this indicates that it responds regardless of where in the visual field an *R* is placed.

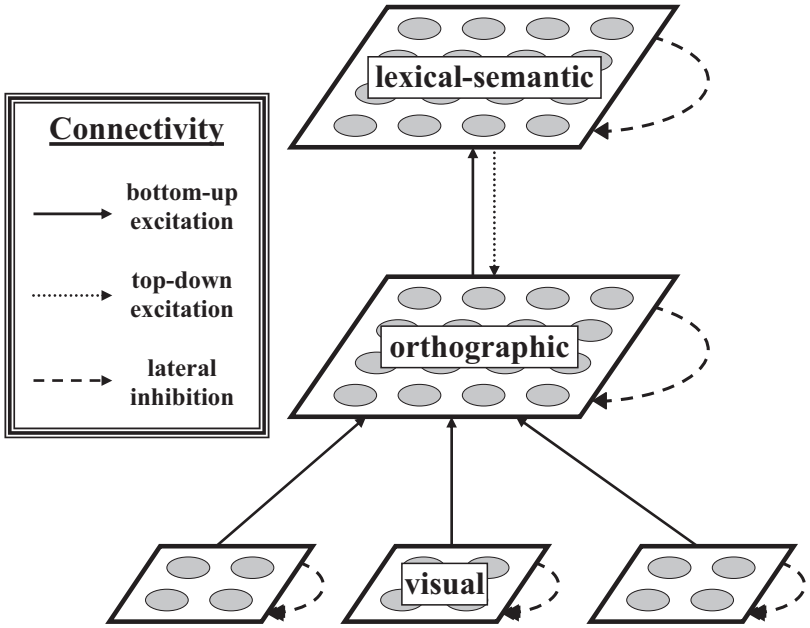


Figure 3. The three-level neural network used for word identification and immediate word priming. The bottom level represents visual features in different spatial locations. The middle and top levels represent the same orthographic or lexical–semantic features regardless of when or where they are presented, resulting in a representation that suffers from spatial and temporal source confusion. Weak top-down excitation from the top level to the middle level gives the model attractor dynamics. Connections between levels are excitatory and specific to learned items. Connections within a level are all-to-all inhibition, representing the actions of inhibitory interneurons. All connections habituate as a function of recent activity due to synaptic depression. From "Persistence and Accommodation in Short-Term Priming and Other Procedural Paradigms: Temporal Segregation Through Synaptic Depression" by D. E. Huber and R. C. O'Reilly, 2003, *Cognitive Science: A Multidisciplinary Journal*, *27*, p. 418. Copyright 2003 by the Cognitive Science Society. Reprinted with permission.

are the same as those contained in O'Reilly and Munakata's (2000) LEABRA (local, error-driven, and associative biologically realistic algorithm) framework, which has been successfully applied to many areas of cognition. The bottom level represents low-level visual information such as simple line segments in particular visual locations. The middle level represents orthographic information regardless of visual location (e.g., a *T* at the beginning of a word regardless of where that word is seen), and the top level represents lexical–semantic information (e.g., a cell representing some aspect of education that is elicited when viewing the word *teach*). With this representation, the same word presented at two different locations or two different times contacts the same units in Levels 2 and 3, resulting in spatial and temporal source confusion. Weak excitatory feedback from the lexical–semantic level to the orthographic level gives the model attractor dynamics (i.e., the ability to focus on the correct identification through resonance), predicting that semantically similar primes may support enhanced orthographic identification of an orthographically dissimilar target (e.g., *doctor→NURSE*). This prediction is tested in Experiment 3, as contrasted with Experiment 2, which used repetition priming. Nonspecific inhibition within each level limits the amount of total activation, which produces masking effects at each level. For instance, if a target word is briefly presented to the model but not masked, then activation at the visual level will linger for hundreds of milliseconds and only slowly dissipate because of leak currents (i.e., the visual icon slowly fades). But, if instead a mask is presented in the same visual location following the target, then the visual representation of the mask inhibits the activation for the visual representation of the word, and this inhibition serves to more rapidly eliminate the word's visual activation. If the mask is a pattern mask, this effect only occurs for the lowest level of the model. But if the mask is itself a string of letters or even another word, then the effect of this inhibition additionally occurs at higher levels, serving to more rapidly eliminate the target word's orthographic and lexical–semantic activations. Such levels-of-masking phenomena are tested in Experiment 1.

Using the architecture shown in Figure 3, Huber and O'Reilly (2003) implemented synaptic depression for each simulated neuron. Additionally, it was assumed that each level of processing accrues information at different rates and, more specifically, that higher levels integrate information more slowly. Because higher levels activate more slowly, this means that higher levels also habituate more slowly than lower levels. For instance, contrast the rapidity of perceptual aftereffects for Figure 2 with the phenomena of semantic satiation (e.g., saying "dog, dog, dog" until the word loses meaning), which takes many seconds or minutes to accrue (Smith & Klein, 1990). The prediction of different rates of processing for the different levels is tested in Experiment 1 by comparing primes that are pattern masks with primes that are words.

Simulating Behavior in Threshold Identification

Critical to understanding the behavior of the model is the decision rule. Because the task presents words at the perceptual threshold, it is assumed that identification is based on partial information. The perceptual fluency decision rule accesses this partial information by selecting the choice word that is identified more quickly. Residual activation from the briefly presented target gives the target word's representation greater fluency than a foil word (i.e., because the target is already slightly active, it more quickly achieves maximum activation). This is the basis of accurate responding when the target word is not explicitly and fully identified. However, because of source confusion, primed foils also receive greater fluency. In sum, there is a greater tendency to choose a target that has been primed (the benefits of priming) as well as a greater tendency to choose a foil that has been primed (the costs of priming). Considering that the same prime presented for a specific duration can produce either benefits or costs depending on the manner of testing, it is the difference in performance between priming the correct answer versus priming the incorrect answer that is the most robust and reliable measure of the magnitude and direction of priming. If the difference between these conditions is positive, the pattern is labeled *positive priming*, but if it is negative, the pattern is labeled *negative priming*.

In what follows, the transition from positive to negative priming is demonstrated by comparing model behavior following a short versus a long duration prime, with both lingering activation (source confusion) and lingering habituation fully enabled but producing differential effects purely as a function of prime duration. The dominating effect of source confusion following brief primes is portrayed in Figure 4, which shows model behavior separately for the representation of the target word (top panels) and the representation of the foil word (bottom panels) for a condition of priming the target with a 50-ms prime (Figure 4A) or priming the foil with a 50-ms prime (Figure 4B) prior to the presentation of the briefly flashed target in the center of the screen (the target was also presented for 50 ms, followed by a 450-ms pattern mask). After the 450-ms pattern-mask presentation, the two choice words appeared, and whichever word achieved its maximum lexical–semantic activation first was selected, as indicated by the time of the vertical gray bars. Because these were deterministic simulations, the difference between the time to peak activation for the target as compared with the foil was mapped to accuracy through the logistic function. The graphs show above-threshold presynaptic activity for the visual, orthographic, and lexical–semantic representations, with these values multiplied by the available nondepleted synaptic resources. Thus, activation in the figure is the output of the cell, which includes the effect of habituation (although habituation does not play much of a role for these brief presentations). The target and prime responses are peaked because the words are immediately masked, and so activation is directly dampened by inhibition from the mask. As can be seen in Figure 4A, priming the target produced two separate visual responses that matched the target: one for the prime and one for the target, even though the repeated target word was presented in different loca-

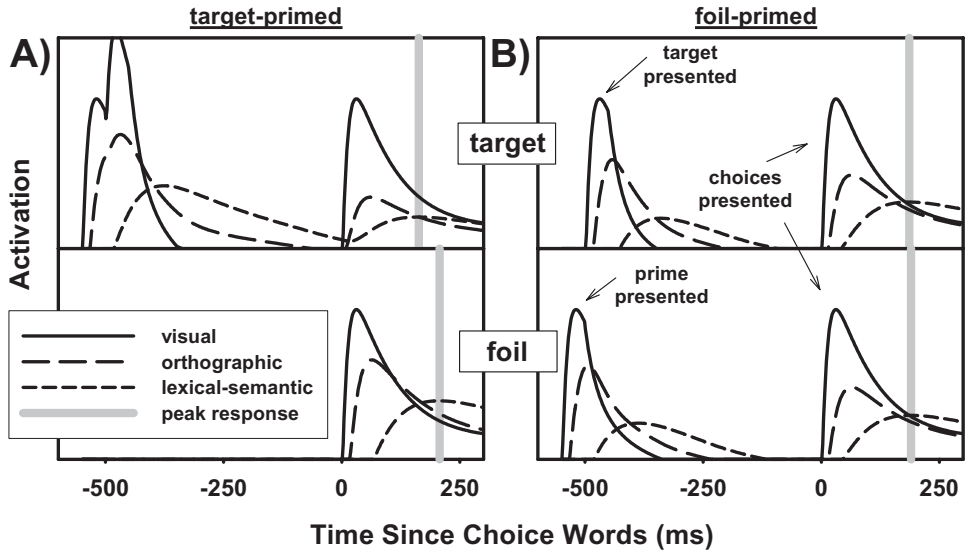


Figure 4. Simulated model behavior that demonstrates positive priming following brief primes due to lingering activation. The curves show activation, which is output from a cell, reflecting the combination of presynaptic activity and the level of habituation through depletion. Portrayed is repetition priming the target with a 50-ms prime (A) or repetition priming the foil with a 50-ms prime (B). In each set of graphs, the top graph shows activation for representations unique to the target word and the bottom graph shows activation for representations unique to the foil word. Each of the three levels shown in Figure 3 is portrayed as a separate curve on the same graph. Five hundred ms after the briefly flashed target, the choice words were presented and a choice was made on the basis of whichever lexical–semantic activation achieved its peak response first, with the timing of this peak response indicated by the vertical gray line. As seen in the foil's representation in B, priming the foil speeded the response to the foil at the time of the choice words, resulting in near chance performance (i.e., very little difference between the two gray lines). Because below-threshold activation exists as well, these graphs do not portray the full extent of residual activation.

tions (there were different pools of visual units for the target at each location, although these are summed together in the figure for convenience). Source confusion between prime and target boosted target activation at all levels, which, in turn, allowed the target choice word to achieve an earlier peak activation than the foil choice word. In contrast, Figure 4B shows a situation in which presenting a prime that matched the foil produced nearly as much residual activation for the foil as residual activation for the target (it may appear that the residual activation completely faded before the choice words were presented, but keep in mind that the figure only shows above-threshold activation). In this foil-primed situation, the residual activation for the primed foil resulted in nearly identical times to peak response for both choice words (i.e., nearly identical perceptual fluency), and, thus, performance was close to chance.

Habituation through synaptic depression serves to counteract and even reverse this preference for primed words for two reasons: First, it reduces the amount of residual activation for primes that have been viewed longer, thereby reducing the preference for primed words; second, lingering depletion for the resources of recently active synapses (i.e., primed synapses) results in sluggish responses when a primed word reappears as a choice alternative. This sluggishness, or disfluency, allows for a full reversal in the direction of preference, resulting in a preference against whichever word was primed. The effect of habituation following long duration primes is highlighted in Figure 5, which compares a condition that primed the target with a 50-ms prime (this is the same as the target portion of Figure 4A) versus one that primed the target with a 500-ms prime. As seen in the lower panel of the figure, even though the prime was continually presented from −1,000 ms to −500 ms, the response to the prime reached a peak and then fell off because of habituation (i.e., unlike in Figure 4, this response falls to a low level even though no masks or other stimuli are presented). When the target flash was presented at −500 ms, the system hardly responded at higher levels because lingering depletion resulted in sluggish responses. To see this, it is important to examine the orthographic and lexical–semantic responses, which achieve approximately half the equivalent response, as seen in the upper panel of Figure 5 (i.e., a cognitive aftereffect). In contrast, the visual response appears normal (i.e., lacking a visual aftereffect). This is because the visual response is not habituated at all considering that the target flash appears in a different screen location than the prime. However, regardless of location, the target corresponds to same orthography and same meaning as the prime, and so these higher level representations are habituated and do not respond readily to the target. In conditions that primed the foil choice word with a long duration prime (not shown), there was a similar sluggish lack of response to the foil's representation at the time of the choice words, which made it easier to select the target.

In summary, because threshold identification is achieved through fluency (i.e., choosing the word that is identified more rapidly), briefly primed words are chosen because residual activation boosts their fluency. Discounting (through habituation) reduces this extra activation and also serves to promote disfluency through sluggish responding due to lingering depletion. The specific assumption that choices are made on the basis of fluency serves as the basis of additional predictions regarding reaction times. More specifically, trials for which the primed word is chosen should reveal faster responding following brief primes but slower responding following long duration primes. These predic-

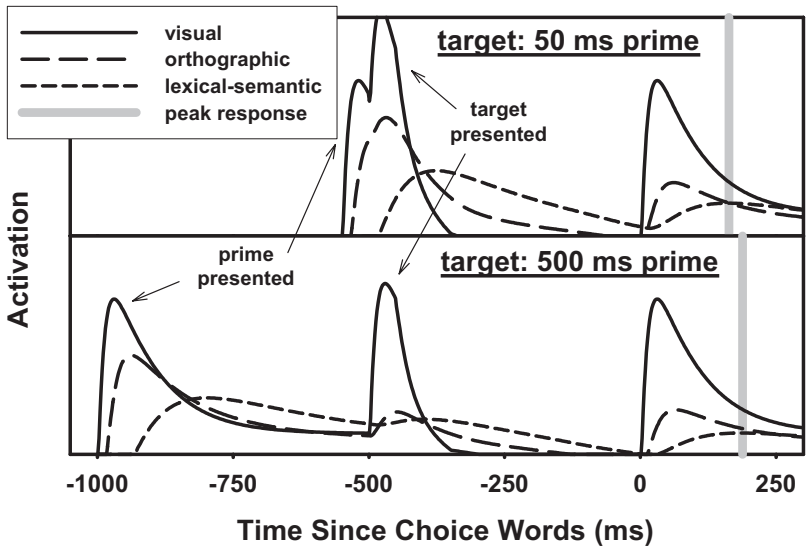


Figure 5. Simulated model behavior that demonstrates negative priming following long duration primes due to habituation. Portrayed is repetition priming the target with a 50-ms prime versus priming the target with a 500-ms prime. In both cases, only the activation of representations unique to the target is shown. After 500 ms of constant exposure to a prime that matched the target (bottom panel), the target's orthographic and lexical–semantic representations were habituated to such an extent that activation in response to the briefly flashed target at −500 ms was approximately half the amount seen in the upper panel. This results in poor performance as indicated by the slower time to achieve peak response to the target's presentation as a choice word. Note that this habituation effect does not hold true for the visual level, which is appropriate considering that the target appears in a different screen location and so the spatially specific visual representation of the target is unaffected.

tions were confirmed both in terms of mean reaction time for correct and error trials and in terms of specific changes in the shapes of the reaction time distributions (Huber & Cousineau, 2004), thus providing converging evidence for this perceptual fluency decision rule.

Overview of Current Experiments

Huber and O'Reilly (2003) developed the neural habituation model prior to collecting these data, and the original motivation for the model was as a response to cognitive aftereffects that appear in many different literatures, including the perception identification results that served as the basis of the ROUSE model. Besides deriving a version of synaptic depression appropriate to real-valued activation rather than spiking neurons, Huber and O'Reilly fit the neural habituation model to a preliminary form of the experiments reported below. However, the details and statistics of these experiments were not presented, and several of the conditions were not included in that article. Additionally, there was no discussion of these results in regard to priming and other areas of psychological research. Therefore, the current article is the first full report of these results, which provide critical tests of the neural habituation model. Nonetheless, there is a close relationship between this article and the Huber and O'Reilly article, with this article presenting the empirical phenomena while that article focused on the mathematics behind neural habitation. The interested reader is referred to Huber and O'Reilly (2003) for the computational specifics of implementing and fitting the model to the experiments reported below.

As seen in the previous section, the Bayesian model of Huber et al. (2001) and the multinomial model of Ratcliff and McKoon (2001) produce negative priming through discounting, but neither model provides predictions in regard to prime duration. Therefore, I focus on neural habituation, which provides explicit predictions regarding prime duration and discounting. Because habituation exists throughout the brain, a change from positive to negative aftereffects as a function of exposure should exist for a large variety of tasks in which there is an immediate effect of one stimulus on a second stimulus or judgment. These extensions and predictions for other situations beyond immediate word priming are considered more fully in the General Discussion. Because the theory is explicitly formulated in terms of biological mechanisms, it also addresses the timing and strength of activation data from electrophysiology and brain imaging data, which are also considered in the General Discussion.

In this article, I report three experiments that test critical aspects of the Huber and O'Reilly (2003) neural habituation model, with each experiment examining the role of prime duration for different stimuli. The first experiment examines different levels of processing by comparing pattern masks, unrelated words, and repeated words, thereby testing the prediction that lower levels of perceptual responses activate and discount more quickly than higher levels and that inhibition plays a role at each level of processing. The second experiment tests the prediction that the deficit for repetition priming when comparing the both-primed condition with the neither-primed condition is due to perceptual habituation and should emerge and continue to grow in magnitude as positive priming (as measured by the difference between the target-primed condition and the foil-primed condition) is eliminated with increasing prime duration. The third experiment tests the prediction that the perceptual deficits seen with repetition priming should change to perceptual benefits with associative priming because of top-down feedback and that these benefits should be maximal when primes maximally mask the target.

Experiment 1: Forward Masking and the Direction of Priming

Experiment 1 provides critical parametric data by testing five different durations of prime presentation, ranging from 17 ms to 2,000 ms. These prime durations are applied to primes that are pattern masks, unrelated words, repetitions of the incorrect foil word, or repetitions of the correct target word (see Figure 1), thereby probing the dynamics of activation and habituation at different levels of processing. The neural habituation model predicts that the transition from positive to negative priming should be gradual, producing a highly nonlinear performance pattern across these prime durations.

Besides testing whether the transition from positive to negative priming is gradual versus steplike as a function of increasing prime duration, Experiment 1 also tests predictions stemming from the use of inhibition within the neural network. To explicate these predictions, I differentiate between priming, which only exists when there is some degree of similarity between prime and target (e.g., repetition priming in this case), and masking, which causes changes in performance even with dissimilar or unrelated visual objects. Thus, changes in the unprimed baseline conditions indicate masking effects. Furthermore, masking can be backward (presentation of visual objects after the target) or forward (presentation of visual objects before the target), and masking can differ as a function of the kind of unrelated visual object that is presented. In summary, presenting unrelated visual objects either before or after the target can make it difficult to see the target, and, furthermore, the exact nature of the unrelated visual objects can serve to modulate the magnitude of the deficit. For instance, an unrelated pattern of lines (a pattern mask) presented before or after the target can have a different effect on performance as compared with an unrelated word (a word mask) presented before or after the target.

In the habituation model, inhibition produces visual masking at the lowest level in the model, and this effect occurs both for masks that precede targets (forward masking) and for masks that follow targets (backward masking). In the case of backward masking, inhibition from the mask serves to more quickly dampen any lingering activation from the target, thus eliminating iconic persistence. In the case of forward masking, lingering activation from the mask inhibits the ability of the target to rapidly and fully activate. Furthermore, these effects can exist at multiple levels of processing, depending on the nature of the mask. Because higher levels of the model process information more slowly, the key prediction is that the strength and timing of masking effects will depend on the type of mask. For instance, pattern masks only mask at the lowest level, whereas word masks produce interference at all levels and should mask more strongly and take longer to achieve their maximal effect. Because the final pattern mask is the same in all conditions (see Figure 1), forward masking serves as the test bed for these predictions.

In general, habituation is expected to produce U-shaped forward masking. Increasing the duration of a very brief mask results in a mask that achieves greater activation and, thus, greater forward

masking effectiveness. However, increasing the duration for a mask that has already been viewed for a while results in a mask that achieves greater habituation (i.e., lower activation) and, thus, lesser forward masking effectiveness. Furthermore, because higher levels activate and habituate more slowly, the timing for the bottom of the U is predicted to vary as a function of the type of mask. For instance, a pattern mask should produce its worst performance for a forward-mask duration that is briefer than that of an unrelated-word forward mask. This is because lexical–semantic processing is slower than visual or orthographic processing and so it takes longer for the activation of the unrelated word mask to habituate. Beyond this levels-of-masking prediction, there should also be a close correspondence between the U-shaped masking effect and the change from positive to negative priming because both effects are due to habituation.

Method

Participants. Sixty-two undergraduates participated, receiving introductory psychology course credit for their participation. All participants were native English speakers with normal or corrected-to-normal vision.

Equipment. Stimulus materials were displayed on PC monitors running real-time MS-DOS (i.e., not running under Windows), with presentation times synchronized to the vertical refresh. The refresh rate was 60 Hz, providing display increments of 16.67 ms. All responses were collected through the computer keyboard.

Materials. One thousand 5-letter words were used for all presentations. These words had a minimum written language frequency of 4 as defined and measured by Kučera and Francis (1967). Randomly generated letterlike pattern masks were used to avoid pattern-mask habituation (see Figure 1 for examples of the pattern masks). All words were displayed in capitalized Times Roman 22-point font. Stimuli were gray against a black background.

Procedure. All variables were within subject and repetition priming was the only type of priming used in this experiment. The basic design used the following two variables: priming condition, with four levels (pattern mask, unrelated word, target primed, and foil primed), and prime duration, with five levels (17, 50, 150, 400, and 2,000 ms). Word selection was random and different for each participant, occurring without replacement such that a given word appeared only once within the experiment, thus avoiding contamination from long-term repetition priming. For the target-primed or foil-primed conditions, the target or foil choice alternative was a repetition of the same prime word (see Figure 1), requiring two different words for a trial of these conditions (as does the pattern-mask condition), whereas unrelated-word trials required three different words per trial.

As seen in Figure 1, the prime word or pattern mask was presented simultaneously in two different locations, both above and below the midline, separated by one pixel. Each trial began with a central fixation point presented for 500 ms, followed by either the prime or five dashed lines along the midline. The duration of the dashed lines was chosen such that the total duration of the dashed lines plus the prime was 2,000 ms (i.e., no dashed lines were presented in the 2,000-ms prime duration conditions). This was done so that participants knew when to expect the briefly flashed target even though prime durations were randomly tested. During the two-alternative forced choice (2-AFC), the target and foil were presented on the left and right of midpoint. The left–right position of the target choice word was fully counterbalanced.

The experiment began with 6 practice trials during which the target duration was set at 133 ms. This was followed by 6 additional practice trials with the target duration reduced to 67 ms. Practice was followed by a block of 80 trials using target durations of 17, 33, 50, or 67 ms, with 20 trials per target duration randomly intermixed across the 80 trials. The purpose of these trials was to determine each participant's perceptual threshold, and participants were informed of this procedure. During both sets of practice trials and the 80 threshold trials, the unrelated-word condition was used and the prime duration was fixed at 400 ms. At the end of the 80 threshold trials, the target duration that placed performance closest to 80% correct was selected and used for the remainder of the experiment. As is typical in perceptual identification experiments, there were large individual differences in the duration needed to obtain the perceptual threshold, with one participant requiring the 17-ms target duration and the other participants even distributed across the other three possible target durations. The average duration was 47.58 ms, with a standard deviation of 13.47 ms.

After the threshold block, there were four blocks of experimental trials with 80 trials per block and mandatory 20-s breaks between blocks. Each block of 80 trials tested the 20 experimental conditions four times, with conditions randomly intermixed across trials and the target as the left choice twice and the target as the right choice twice. After collapsing across left–right position and across the four blocks, this yielded 16 trials per condition per participant in the reported results.

Participants were fully informed that it was just as likely that a primed choice word would be the correct choice as the incorrect choice and that there could be no effective strategy in relation to the primes except to stay focused on the center of the screen and try to see the briefly flashed target. These claims were reinforced through trial-by-trial error feedback. The 2-AFC was nonspeeded and every trial was self-paced, with participants hitting a key once they read the error feedback and again when they were ready for the next trial.

Results and Discussion

Average probability correct values are shown graphically in Figure 6 and numerically in Appendix A. There was a main effect of priming condition, $F(3, 58) = 101.12$, $p < .001$, and a main effect of prime duration, $F(4, 59) = 41.36$, $p < .001$, and these two variables interacted, $F(12, 50) = 21.53$, $p < .001$. Next, I consider the masking effects (upper panel of Figure 6) followed by the repetition priming effects (lower panel of Figure 8), providing specific contrasts designed to test the predictions of the neural habituation model.

Forward masking. The model includes nonspecific inhibition within each level of processing and predicted that masking effects should be stronger for unrelated words than for pattern masks and that these forward-masking effects should reveal a U-shaped pattern due to the buildup of activation that drives the subsequent buildup of habituation. Because a pattern mask only involves visual masking and because higher levels of perception activate and discount more slowly than lower levels, the U-shaped pattern

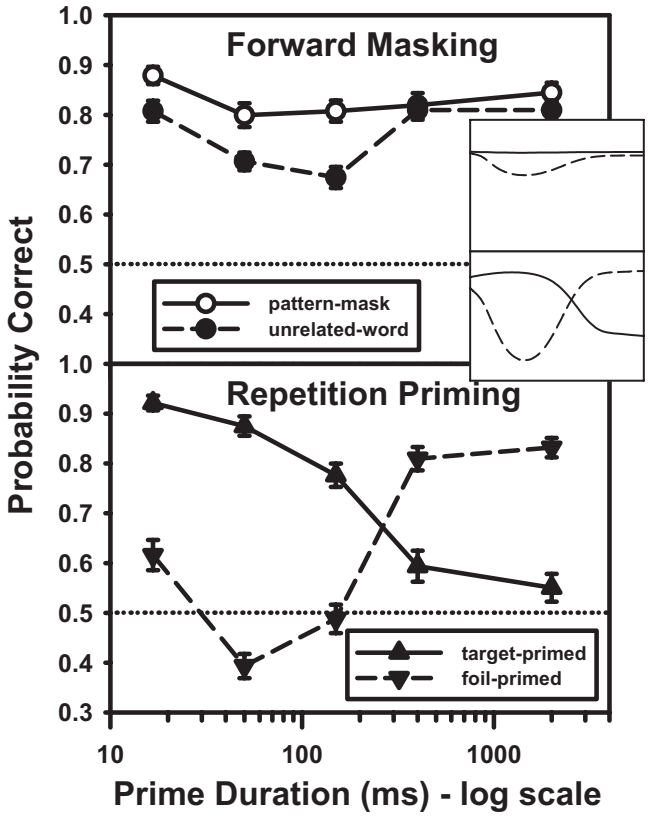


Figure 6. Forced-choice performance in Experiment 1 with the dashed line at .5 showing the chance level. The inset shows the simulated behavior of the neural habituation model developed by Huber and O'Reilly (2003). The upper panel shows accuracy for the pattern-mask and unrelated-word conditions as a function of prime duration on a log scale, demonstrating the forward masking characteristics of these stimuli. As predicted, an unrelated word prime is a more effective mask than a pattern-mask prime, and both types of masks reveal U-shaped patterns as a function of prime duration, with the largest decrement occurring at shorter prime durations for pattern masks. The lower panel shows accuracy for the target-primed and foil-primed conditions. As predicted, there is a clear pattern of costs and benefits, with performance better with target priming and worse with foil priming for brief prime durations, but the opposite is observed for long prime durations, with performance better with foil priming and worse with target priming. The error bars show ±1 standard error of the mean.

should reach its maximum deficit more quickly for pattern masks as compared with words.[3]

Collapsing over prime duration, performance was better in the pattern-mask condition as compared with the unrelated-word condition, $t(61) = 6.96$, $p < .001$, confirming the prediction regarding levels of masking with forward masks. Similar levels of masking effects have been observed for backward masks in visual perceptual word identification (Taylor & Chabot, 1978), but forward masking is not commonly examined for levels of masking effects. Visual inspection of Figure 6 reveals that the lowest level of performance for the pattern-mask condition occurred in the 50-ms condition (U-shaped planned comparison across 17, 50, and 150 ms), $t(61) = 3.04$, $p < .0025$, whereas the lowest level of performance for the unrelated-word condition occurred in the 150-ms condition (U-shaped planned comparison across 50, 150, and 400 ms), $t(61) = 4.79$, $p < .001$, confirming not only that words are more effective forward masks than are patterns but also that words reach their maximum forward masking potential at longer durations.

Beyond its significance to priming phenomena, this U-shaped forward masking effect may relate to experiments examining visible persistence (i.e., the tendency for a visual icon to linger for a few hundred milliseconds beyond the time when a display disappears; Sperling, 1960). A classic paradigm that examines persistence is the "missing dot" task in which simultaneous perception of a first presentation and a subsequent presentation is needed to determine which spatial location does not contain a dot in either display (Hogben & Di Lollo, 1974). In this task, increasing the duration of the first display initially helps performance, presumably because of increased visible persistence of the first display as it gains stimulus energy. However, longer durations of the first display decrease performance, presumably because of habituation to the first display resulting in decreased visible persistence. The latter result with long durations is termed the *inverse duration effect*. For the missing dot paradigm, visible persistence enables performance, but with forward masking effects, visible persistence interferes with the subsequent perception. In other words, both are due to simultaneous perception of the previous and subsequent displays, but in the case of the missing dot paradigm, this simultaneity helps identify what is missing, whereas in an identification paradigm, the simultaneity interferes with accurate perception. This task difference explains why the missing dot paradigm produces n-shaped performance, whereas the masking effects observed in Experiment 1 revealed U-shaped performance as a function of prime duration.

Repetition priming. The neural model predicted that brief prime durations result in a tendency to choose whatever has been primed because of lingering activation and source confusion. These priming effects were expected to help performance in the target-primed condition and harm performance in the foil-primed condition. In contrast, long-duration primes were expected to produce the opposite pattern of results because of habituation and sluggish responding of a primed choice word resulting in a tendency not to choose whatever was excessively primed. This transition from positive to negative priming was expected to progress in a gradual manner rather than the discrete manner that otherwise might be associated with strategic discounting.

As seen in Figure 6, the accuracy differences between the target-primed and foil-primed conditions were large and varied with prime duration. The target-primed condition was significantly better than the foil-primed condition for the 17-, 50-, and 150-ms conditions, in that 60 of 62 participants showed this effect, $t(61) = 14.38$, $p < .001$, but the opposite was true for the 400- and 2,000-ms conditions, in that 52 of 62 participants showed this effect, $t(61) = 7.27$, $p < .001$. For short prime durations, participants were apparently very confused as to which word was the

[3] It is difficult to tell from the inset of Figure 6, but there is indeed a very small deficit for pattern masks presented for 50–10 ms before the target. The model readily produces a more sizable deficit than is apparent, but in fitting the model to these 20 conditions with just six free parameters, the best solution minimized the pattern-masking effect. The chi-square error goodness-of-fit measure resulted in much more pressure to fit the greatly varying target-primed and foil-primed conditions, and, thus, fitting the pattern-mask conditions was to some degree sacrificed.

prime and which word was the target. For instance, at the 50-ms foil-primed condition, performance was significantly below chance, $t(61) = 4.36$, $p < .001$, indicating that responses in this condition were more strongly influenced by the prime than by the briefly flashed target.

The pattern seen in the lower panel of Figure 6 might seem confusing at first if one's naive expectation is for a symmetric crossover pattern between the target-primed and foil-primed condition as a function of prime duration. However, it is important to keep in mind that forward masking plays a role even for these conditions (i.e., the prime also acts as a mask) and so they are best compared with the unrelated-word condition, which revealed a U-shaped pattern as a function of prime duration. As compared with this U-shaped unprimed condition, these repetition priming conditions do indeed produce a symmetric crossover pattern between positive priming for short prime durations and negative priming at longer prime durations.

According to the dynamic neural network account, habituation is the underlying cause of both the upturn in the U-shaped masking pattern for unrelated words as well as the crossover pattern with target priming, and, as predicted, these two phenomena are coupled, with both reaching appreciable levels at a duration falling between 150 and 400 ms. Furthermore, the theory explains why these effects occur around 150 ms for a well-attended visually presented word; neural habituation does not have a set time course, and it is the process of identifying the prime word that drives the habituation. Therefore, the theory predicted a close correspondence between reading speed and the onset of discounting due to habituation.

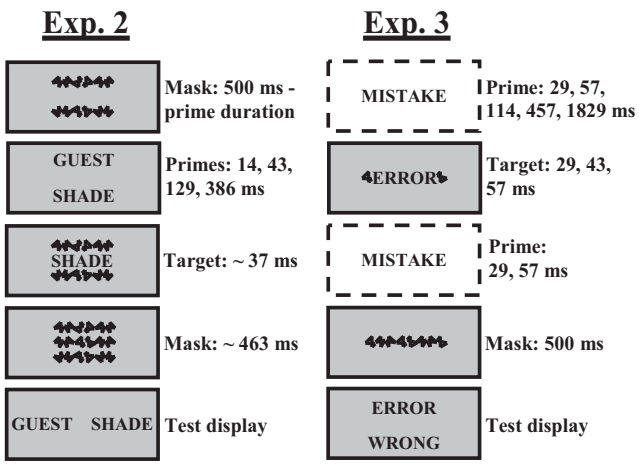


Figure 7. Both-primed conditions for repetition priming (Experiment 2) and associative priming (Experiment 3). Experiment 2 set target duration separately for each individual, whereas Experiment 3 used three specific target durations. In Experiment 3, the prime could appear either before or after the target, as indicated by the dashed boxes. The target and foil words always contained the same number of letters in Experiment 3, but that number could vary from three to seven letters and all primes and targets were presented with sufficient flanking pattern-mask characters to cover seven letter positions. Not shown are the unrelated word (both experiments), target-primed (Experiment 2), foil-primed (Experiment 2), and pattern-mask (Experiment 3) conditions.

Experiment 2: Both-Primed Deficits With Repetition Priming

Experiment 1 revealed a gradual change from positive to negative priming (gradual accrual of habituation), U-shaped forward masking effects (inhibition between any two visual objects), different timing in U-shaped forward masking in comparing pattern masks with unrelated words (higher levels process information more slowly), and, finally, a close correspondence between the onset of negative priming and the duration of maximum masking (habituation is the underlying cause of both effects). As explained next, the neural habituation model differs from the ROUSE model and the multinomial model because it assumes that priming results in real perceptual changes rather than changes in the evaluation of perceptual evidence. This key difference is tested both in Experiment 2 and in Experiment 3.

Experiment 1 included conditions without priming (the masking conditions) as well as repetition priming of one choice alternative to assess whether priming was positive or negative. Forced-choice testing allows for another condition in which both choice words are equally primed. For repetition priming, this both-primed condition requires the presentation of two different prime words simultaneously (see Figure 7 for an example of the both-primed condition for repetition priming, Experiment 2, and the both-primed condition for associative priming, Experiment 3). This condition is unbiased in that any decision strategy in relation to the primes (e.g., "If I can't see the flashed target, then I'll use the strategy of not choosing the primed choice") would apply equally to the choice words and, thus, be incapable of affecting performance. Therefore, this both-primed condition is a critical test of priming effects above and beyond strategic responding.

In the Bayesian ROUSE model, perception is unchanged (i.e., the number of features activated by the target is unaffected by priming), and the observed priming effects are due to a decisional preference for or against primed features. A preference for primed features helps performance in the target-primed condition and hurts performance in the foil-primed condition. Assuming that perceptual features are activated probabilistically by presentation of a prime, there will be trial-by-trial variability in the magnitude of these preference effects, and so the ROUSE model actually predicts deficits when both target and foil are primed. Essentially, it is a signal to noise problem, with the signal coming from the valid features of the flashed target and extra noise features coming from priming. The same is true of the multinomial model, only the noise is all-or-none source confusion and all-or-none discounting from the primes that serves to override perception. In keeping with these claims, a variety of experiments have revealed robust both-primed deficits with repetition priming or orthographic priming (Huber, Shiffrin, Lyle, & Quach, 2002; Huber et al., 2001; Huber, Shiffrin, Quach, & Lyle, 2002; Weidemann et al., 2005).

In contrast to this account of both-primed deficits based on extra variability with priming, the neural dynamics of Huber and O'Reilly (2003) produce both-primed deficits through a true perceptual deficit in which there is less of a perceptual response to the target. More specifically, as habituation accrues with increasing prime duration, the system becomes less able to activate the briefly flashed target word.

The ROUSE model predicts that the both-primed deficit is related to the degree of source confusion (which produces variability) but not discounting, and the multinomial model predicts

that the both-primed deficit is a function of both source confusion and discounting (both of which are variable). In contrast to both of these assumptions, the neural habituation model assumes that the deficit is related to discounting through habituation. Thus, both the ROUSE and multinomial models suppose that the onset of positive priming should correspond to the onset of both-primed deficits (i.e., a prime sufficient to produce priming should also produce both-primed deficits through variability), whereas the habituation model supposes that both-primed deficits will lag behind the onset of positive priming because habituation only accrues in response to activation (i.e., both-primed deficits only arise after sufficient priming). Therefore, Experiment 2 tested very short prime durations to determine whether positive priming accrues in advance of both-primed deficits. Such a finding would provide evidence against the variability explanation supposed in the ROUSE and multinomial models: Only the habituation model allows for positive priming in the absence of both-primed deficits.

Repetition priming of both choice words requires use of two different prime words on every trial. Huber et al. (2001) did this by presenting the two prime words above and below fixation with some space between each prime (Figure 1 shows no separation between the above and below primes, whereas in Experiment 2, one word was presented above and a different word was presented below, with space between the two words). However, this method of presentation results in priming that is not as strong or as fast in its transition between positive and negative priming as compared with the results of Experiment 1 (e.g., Huber, Shiffrin, Quach, & Lyle, 2002). This may be partly due to the more peripheral display of the primes (Weidemann et al., 2005). In addition, the neural habituation model predicts this weakening of priming effects due to the extra inhibition that comes from the presentation of two different prime words. Regardless of the specific underlying causes, it is important to keep in mind that the results of Experiment 2 are not quantitatively comparable to Experiment 1 in terms of prime duration because the former used two prime words that were more peripheral and the latter only used one prime word. In the current situation, it is useful that this mode of presentation slows down the transition from positive to negative priming, because the critical prediction for this experiment was that both-primed deficits will lag behind the accrual of positive priming; testing such a prediction requires the use of several brief prime durations prior to the onset of discounting. Finally, unlike Experiment 1, Experiment 2 presented pattern masks in advance of the primes, which served to further slow down the ease with which primes are identified.

Method

All materials and procedures were identical to those used in Experiment 1 except where otherwise specified. Unlike Experiment 1, every trial of Experiment 2 presented two different prime words that were vertically separated from each other (see Figure 7). This method of presentation was the same as in Experiment 1 of Huber, Shiffrin, Quach, and Lyle (2002), which tested prime durations of 500 and 2,500 ms. However, the current experiment filled in the entire time course of prime duration down to the level of subliminal priming by testing prime durations of 14, 43, 129, and 386 ms.

Participants. Forty-four university undergraduates participated, receiving introductory psychology course credit for their participation. All participants were native English speakers with normal or corrected-to-normal vision.

Materials. All words were displayed in capitalized, Times Roman 12-point font. Stimuli were black against a gray background.

Equipment. The refresh rate was 70 Hz, providing display increments of 14.29 ms. To present the many different rapid displays involved in this experiment, screen resolution was lowered to 320 × 200 pixels, yielding a sufficient number of video pages.

Procedure. The basic design used the following two variables: priming condition, with four levels (unrelated words, both primed, target primed, and foil primed), and prime duration, with four levels (14, 43, 129, and 386 ms). Three experimental blocks of 64 trials each included 4 repetitions of all 16 possible priming conditions, yielding 12 repetitions per condition per participant.

Words and pattern masks were presented in three different horizontal lines of text placed one above the other. The middle line was reserved for the fixation point, the briefly flashed target, and the pattern mask that backward masked the target. The upper and lower lines were reserved for the two prime words and pattern masks that both forward masked and backward masked the primes. Every trial began with a fixation point for 500 ms followed by a blank screen for 500 ms. This was followed by pattern masks at the upper and lower prime locations and then the prime words at these locations (the middle line remaining blank). The combined time of the pattern masks and primes was 500 ms (e.g., a 14-ms prime consisted of 486 ms of pattern masks followed by 14 ms of prime words). In this manner, the duration from the onset of the fixation point to the target flash was kept constant at 1.5 s across all trials. During the briefly flashed target and during presentation of the central pattern mask that replaced the target, pattern masks were presented in the prime word locations to fully mask the primes.

Forced-choice responses were collected on an 8-point scale between the left and right choice words, representing the differential degree of confidence in the chosen word. A staircase method was used to progressively home in on the target duration that produced correct responses 75% of the time. This staircase method allows a wide range of possible threshold target durations to suit the large individual differences. The average target duration was 37 ms, although individual target durations ranged from 14 ms to 128 ms.

Results and Discussion

Average probability correct values are shown graphically in Figure 8 and numerically in Appendix B. There was a main effect of priming condition, $F(3, 41) = 13.75$, $p < .001$, and a main effect of prime duration, $F(3, 41) = 47.19$, $p < .001$, and these two variables interacted, $F(9, 35) = 5.55$, $p < .001$. It is useful to break the four priming conditions into the two that are unbiased (upper panel of Figure 8) versus the two that indicate the direction and strength of preference (lower panel of Figure 8). The former are termed *unbiased* because either the two prime words were equally unrelated to the choice words or the two prime words were repetitions of the two choice words. Therefore, explicit identification and use of the prime words would not differentially indicate a particular choice. The other two priming conditions are termed *preference conditions* because just one of the two choice words was primed, and so these conditions measure whether the direction

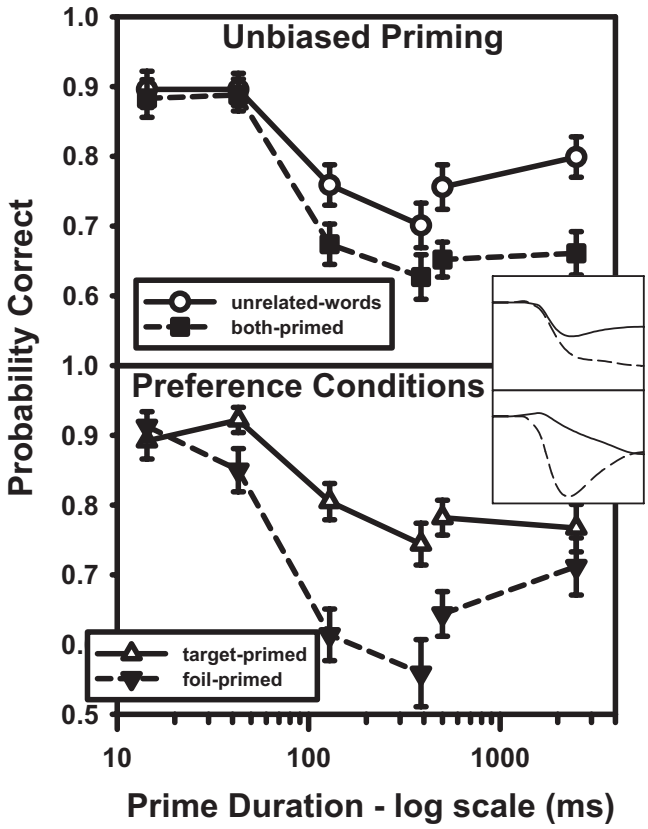


Figure 8. Forced-choice performance in Experiment 2, which tested repetition priming. The inset shows the simulated behavior of the neural habituation model reported by Huber and O'Reilly (2003). This experiment included the same four priming conditions as the short (500-ms) and long (2,500-ms) conditions of Huber, Shiffrin, Quach, and Lyle (2002), which are also shown in the graph. Unlike that experiment, Experiment 2 tested shorter prime durations (14, 34, 129, and 386 ms). Unlike Experiment 1, two different prime words were presented on every trial rather than the same prime word being shown twice. Furthermore, the two prime words were more vertically separated than in Experiment 1, with a blank line in the center of the screen between the top and bottom prime words, and the primes were forward and backward masked by pattern masks. These changes served to slow down the accrual and subsequent elimination of priming as a function of prime duration. As predicted, the both-primed deficits lag behind positive priming and then gradually increase in magnitude as the preference for primed words decreases. The error bars show ±1 standard error of the mean.

of priming was for or against primed words. These are nominally the same as the repetition conditions of Experiment 1, except that (a) two different prime words were presented on every trial, (b) primes were more peripheral, and (c) primes were both forward and backward masked by pattern masks. All three of these methodological changes serve to reduce the magnitude priming and slow down the transition from positive to negative priming.

The neural habituation model predicted that both-primed deficits (the difference between the unrelated-word and both-primed conditions) should increase (or remain unchanged) with increasing prime duration but that the preference (the difference between the target-primed and foil-primed conditions) should at first increase and then decrease. For the new data of Experiment 2 that tested the 14-, 34-, 129-, and 386-ms prime durations, a linear contrast across the four prime durations revealed increasing both-primed deficits, $t(43) = 2.67$, $p < .01$, combined with increasing positive preference, $t(43) = 3.94$, $p < .001$, as a function of increasing prime duration. For the previously published data of Huber, Shiffrin, Quach, and Lyle (2002) that tested the 500- and 2,500-ms prime durations, a linear contrast across the two prime durations revealed no significant change in both-primed deficits, $t(32) = 0.68$, $p = .50$ (although there was a trend for increasing both-primed deficits), combined with decreasing positive preference, $t(32) = 1.99$, $p < .05$, as a function of increasing prime duration. If instead the data had revealed decreasing both-primed deficits at the same time positive preference decreased, this would have constituted a falsification of the habituation model, which supposes both-primed deficits are perceptual in nature and due to discounting.

Besides this prediction that both-primed deficits should only increase (or remain the same) with increasing prime duration, the neural habituation model also predicted that the both-primed deficits should lag behind the emergence of a positive preference. This is because the both-primed deficits are due to the loss of target perception that happens with neural habituation (discounting), whereas the positive preference is due to the integration of activation from the primes (source confusion). This lag arises in part from the threshold of activation in the simulated neurons. A sufficiently brief prime presentation can activate the prime's representation to a subthreshold level, which enhances the response to a primed target or primed foil without triggering habituation; only above-threshold activity drives neural habituation, and so a subthreshold prime presentation will not produce any noticeable decrement in target perception. This pattern was seen at the 34-ms prime duration. Paired comparisons between the target-primed and foil-primed conditions revealed a positive preference at this duration, $t(43) = 3.39$, $p < .001$, despite the absence of both-primed deficits, $t(43) = 0.38$, $p = .71$. This particular result at 34 ms is in disagreement with the variability explanation of both-primed deficits in the ROUSE and multinomial models—the observation of sufficient source confusion to produce a positive preference should have also produced both-primed deficits according to these models because of variability in the amount of source confusion.

Experiment 3: Both-Primed Benefits With Associative Priming

Experiment 2 provided evidence in support of a perceptual basis of both-primed deficits due to lack of perceptual response to the target word (i.e., habituation) rather than a deficit based on variability. However, this support was somewhat indirect and relied on null hypothesis testing (i.e., a lack of both-primed deficit at a prime duration that also produced positive priming). Providing more definitive evidence in favor of this perceptual account of both-primed effects, Experiment 3 tested a key qualitative difference between the models. Unlike the other two models, the neural habituation model supposes that associative-semantic priming produces both-primed benefits (i.e., actual perceptual enhancement). The other models cannot produce such both-primed benefits. In support of the neural habituation model, past work found small but reliable both-primed benefits with associative priming (Huber et al., 2001; Masson & Borowsky, 1998). Experiment 3 replicated these both-primed perceptual benefits with associative priming

while additionally examining how these effects relate to prime duration and forward masking.

In the habituation model, both-primed benefits arise from top-down support from semantically similar primes (perceptual support). In contrast, both-primed deficits are due to habituation of orthography (lack of perceptual response due to habituation). In the case of repetition priming, top-down support exists but is insufficient to offset the deficits due to orthographic habituation. However, semantic priming does not induce orthographic habituation, and, thus, semantic priming provides the best opportunity to observe perceptual benefits. Top-down semantic support produces below-threshold activation (e.g., the word *doctor* makes you think of the letters for *nurse*, but not so much that you hallucinate the letters for *nurse*), and, thus, there is some orthographic benefit for semantically related words but no deficit due to habituation.

The key prediction was that the both-primed benefits should be largest at the prime duration that also produced the largest forward masking effect (poorest performance in the unrelated-word condition). In other words, forward masking was expected to again reveal a U-shaped pattern as a function of prime duration, and, furthermore, it was predicted that the prime duration at the bottom of the U would also reveal maximal both-primed benefits. Thus, a tight linkage was expected between masking and priming as a function of prime duration. This is because lingering semantic activation from the prime is both the source of semantic forward masking effects as well as the source of semantic both-primed benefits. Therefore, the prime duration that produces the most source confusion at the semantic level is the prime duration that will produce the most competition for an unrelated target but, at the same time, the most top-down support for the orthography of a semantically related target.

In the model, activation accrues somewhat slowly and different presentations blend together (i.e., source confusion), with concurrent activation from subsequent presentations. Therefore, the model predicted that both-primed benefits can arise in response to primes that precede the target but also that both-primed benefits can arise from primes that follow the target.[4] In other words, top-down semantic support from a prior prime can aid subsequent orthographic processing, but it is also possible that subsequent semantic processing can aid as yet unresolved orthographic processing of the target when the prime follows the target. Finally, providing further validation for the levels of forward masking effects seen in Experiment 1, Experiment 3 included conditions that presented pattern masks at different durations rather than prime words.

To produce larger priming effects, Experiment 3 used a single prime word, similar to Experiment 1. Because no condition included repetition priming, it was not necessary to present the prime word in two spatial locations (for Experiment 1, two identical primes were used to allow demarcation of the target as the only word in the center). Therefore, the prime word was presented in the same position as the flashed target word (see Figure 7). Furthermore, because the prime appeared after the target in some conditions, it was not known in advance whether the identity of the first or the second word would be tested in the forced-choice decision. In this design, the strategy to choose whichever word was seen (regardless of when it was seen) would nevertheless support accurate responding because no condition used repetition priming. This uncertainty for which word was the target in the sequence of events ensured a high level of attention to both words presented within the trial sequence. An associative-semantic both-primed condition was constructed by finding two different semantic associates for every prime word. For instance, the prime *MISTAKE* might be followed by the target *WRONG* and then a choice between *WRONG* and *ERROR.*

Method

All materials and procedures were identical to those used in Experiment 2 except where otherwise specified.

Participants. Ninety-four university undergraduates participated, receiving introductory psychology course credit for their participation.

Materials. Two hundred sixty four prime/Associate 1/Associate 2 (e.g., *MISTAKE*/*WRONG*/*ERROR*) triples were created. Primes and associates ranged from 3 to 7 letters in length, with prime length and associate letter length often differing, but both associates were constrained to the same number of letters (otherwise the forced choice could be achieved through word length). Either associate could appear as a target or foil, thereby counterbalancing association strength across participants. The average forward association strength from prime to associate was .145 (D. L. Nelson, McEvoy, & Schreiber, 1994). Twelve-point Geneva font was used for all displays.

Procedure. There were 66 conditions in the experiment, 63 of these resulting from 3 target durations by 3 types of primes by 7 prime durations (see Figure 7). In addition, during three target durations, no prime was presented (0-ms condition). Unlike all other experiments, there was no block of trials to determine the individual target duration threshold. Instead, the same three target durations were used throughout the experiment and the reported results averaged across these three target durations. Because all words appeared in uppercase and targets could appear either after or prior to primes, participants had no way of identifying which word was the target until the forced-choice options appeared. In other words, it was necessary to wait for the choice words and then pick the one that was previously seen, regardless of whether that word appeared first or second. Each of the conditions occurred four times during the experiment (12 data points per condition per subject after collapsing across the three target durations). Target durations were 29, 43, and 57 ms. The three types of primes were pattern mask, unrelated word (created by re-pairing prime and associates), and both primed (intact prime–associate presentation). Target and foil were always the two associates of a single prime (even on unrelated-word and pattern-mask trials). In other words, participants were always given a choice between two words that were both associates of a single word (even if that word was not shown). The seven prime durations consisted of five conditions in which the prime or mask preceded the target and two conditions in which the prime or mask followed the target. The five preceding times were 29, 57, 114, 457, and 1,829 ms. The two following times were 29 or 57 ms.

Presentation order was 500-ms fixation, 500-ms blank screen, prime (if the prime preceded), target, prime (if the prime fol-

[4] Presenting primes after the target may seem to violate the notion of priming the pump, but considering the sluggishness of perceptual processing, presenting primes either immediately before or immediately after the target was expected to produce a substantial amount of concurrent prime and target activation, thus producing sizable priming.

lowed), 500-ms pattern mask, 500-ms blank screen, and then forced choice until response. Because primes and targets were of varying lengths from three to seven letters, it was necessary to flank primes and targets with additional pattern-mask characters for words that were fewer than seven letters in length and it was necessary to increase the number of pattern-mask characters in the final mask to seven. This ensured that any prime or target was fully masked by subsequent presentations. Except for the forced choice, all items appeared in the same location in the center of the screen. Unlike previous experiments, the forced choice consisted of one word above and the other below the central line. Furthermore, the final 500-ms pattern mask was not reduced by the target duration. However, if the prime appeared after the target, the duration of the final pattern mask was reduced by the duration of the prime. This served to equate the time between offset of the target and presentation of the forced choice.

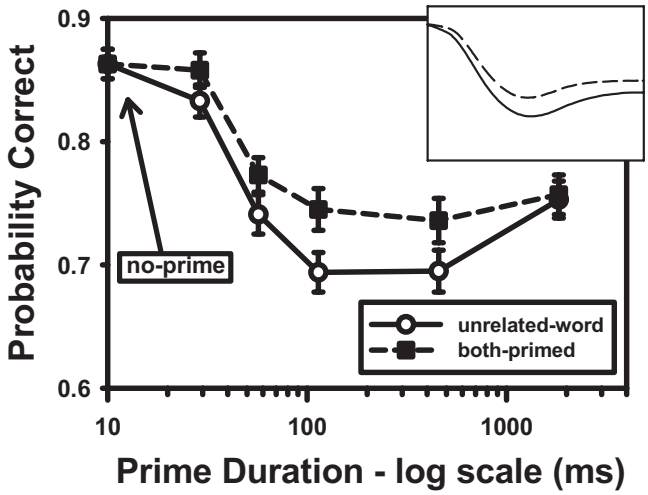


Figure 9. Forced-choice performance in Experiment 3, which tested associative priming (e.g., *MISTAKE* priming *WRONG* and *ERROR* as choices). The inset shows the simulated behavior of the neural habituation model reported by Huber and O'Reilly (2003). Experiment 3 additionally included pattern-mask conditions as well as conditions that presented the prime after the target (shown in Appendix C). Similar to Experiment 1, only a single prime word was used on every trial. Unlike Experiment 1, the prime was presented in just one location in the center of the screen. As predicted, associative priming produced both-primed benefits, and these benefits achieved their largest value at the prime durations that produced the largest degree of forward masking. The error bars show ±1 standard error of the mean.

Results and Discussion

Unlike Experiments 1 and 2, which determined the target threshold duration separately for each participant, Experiment 3 used the same three target durations for all participants. The hope was that at least one of these durations would provide for threshold performance. Consistent with an increasing psychometric function with increasing target duration, there was a main effect of target duration, $F(2, 92) = 320.17$, $p < .001$, but no interaction between target duration and priming condition, $F(2, 92) = 1.66$, $p = 20$, and no interaction between target duration and prime duration, $F(12, 82) = 1.29$, $p = .24$, and so the results were collapsed over target duration to simplify the analyses, reducing the original 66 conditions to just 22 conditions.

Average probability correct values for the unrelated-word and both-primed conditions in which the prime preceded the target are shown graphically in Figure 9 along with the modeling results. In addition, probability correct for all 22 conditions is reported numerically in Appendix C. A 3×7 analysis of variance for priming condition and prime duration (including all conditions except the single no-prime condition) revealed a main effect of prime duration, $F(6, 88) = 33.55$, $p < .001$, a main effect of prime condition, $F(2, 92) = 12.30$, $p < .001$, and an interaction between priming condition and prime duration, $F(12, 82) = 16.94$, $p < .001$. Next, specific contrasts are reported to test qualitative model predictions.

As measured by separate *t* tests, the both-primed advantage over the unrelated-word condition was significant at the .05 level or better for all prime durations except the longest prime preceding target duration, 1,829 ms, and the shortest prime following target duration, 29 ms. This demonstrates that an associatively related word can produce a real perceptual benefit and, furthermore, that this benefit can occur if the prime is presented either before or after the briefly flashed target. Both the ROUSE model and the multinomial model are incapable of producing any perceptual benefits. In contrast, the neural habituation model produces these effects because of top-down support and also explains why (a) these effects exist for associative-semantic priming but not orthographic or repetition priming, considering that the latter additionally includes perceptual deficits with orthographic habituation; (b) these effects exist for priming both before and after the target, considering that perceptual activation is blended over time and across presentations; and (c) the largest perceptual benefits are seen at durations that produce the largest forward masking effects in the unrelated-word condition, considering that lingering prime activation at the semantic level is the source of both phenomena.

As with Experiment 1, it was expected that pattern masks and unrelated words presented before the target would produce U-shaped forward-masking effects. Furthermore, it was predicted that pattern masks would achieve their lowest performance level at shorter prime durations than would unrelated words considering that semantic activation and semantic habituation accrue more slowly as compared with visual features. The data confirm these predictions, and pattern-mask primes (see Appendix C) reached their lowest performance level at 57-ms prime duration whereas unrelated-word primes produced their lowest performance level at 114-ms prime duration. These claims were statistically addressed through piecewise linear contrast weights on the forward masking conditions of +1, −1, −.5, 0, +.5, for the pattern-mask condition, $t(93) = 6.59$, $p < .001$, and +1, 0, −1, 0, +1, for the unrelated-word condition, $t(93) = 6.88$, $p < .001$. To test the prediction that the largest both-primed benefit occurs at the same target duration as the largest forward masking deficit for the unrelated-word condition, the same +1, 0, −1, 0, +1, contrast was used on the differences between the both-primed and unrelated-word conditions, $t(93) = 1.84$, $p < .05$. However, unlike Experiment 1, the magnitude of pattern masking was overall no different than the unrelated-word condition for these five durations that presented the prime before the target, $t(93) = 1.82$, $p = .072$. Unlike Experiment 1, which used Times Roman 22-point font, Experiment 3 used Geneva 12-point font (because of lower screen resolution), and this difference may have introduced differential masking effectiveness of the pattern masks as compared with letters

between the two experiments. It is also important to note that differential masking effectiveness cannot be directly compared between Experiments 1 and 3 because Experiment 1 set target duration individually for each participant to obtain threshold performance, whereas Experiment 3 used fixed target durations.

General Discussion

Each of the presented experiments tested key predictions of the neural habituation priming model proposed by Huber and O'Reilly (2003). Assuming that habituation exists at many levels of processing, including visual, orthographic, and lexical–semantic representations, the reported results were interpreted as cognitive aftereffects, similar to perceptual aftereffects that produce a negative image after overexposure. The term *cognitive aftereffect* is used to broadly refer to habituation within any type of identification process, whether that be for detection of simple visual line segments, detection of letters in specific positions within a word, or detection of specific forms of meaning implied by a word. Such aftereffects are hypothesized to arise from habituation, which may serve as an adaptive mechanism for reducing source confusion between subsequent presentations. This boosts performance when the prime and target are unrelated because primed words are eliminated from the list of possible perceptions (the foil-primed condition). However, there is a cost to habituation, as primed targets are not as easily perceived (the target-primed condition). All three experiments used forced-choice testing of words presented at threshold to measure perceptual effects while minimizing strategic bias. In addition, strategies were discouraged by using trial-by-trial accuracy feedback and a mix of conditions that was just as likely to prime the wrong answer as the correct answer. Because there could be no effective strategy that used the primes, this allowed the examination of priming effects even with highly salient primes. Unlike previous work with the paradigm, all three experiments used parametric variations in prime duration to test the dynamics of priming.

Experiment 1 examined repetition priming with single primes and found large priming effects as a function of prime duration. There was a tendency to choose primed words that first increased and then gradually decreased as a function of increasing prime duration. For even longer prime durations, the tendency switched to a preference against choosing primed words (i.e., negative priming). In addition, there was an important correspondence between these priming effects and forward-masking effects such that the switch from positive to negative priming corresponded to a lessening of the forward masking from an unrelated prime word. According to the neural habituation model, priming and forward masking are both due to lingering activation (source confusion) from prime presentations, but this lingering activation is reduced by habituation (discounting) following long viewing durations. Assuming that threshold performance is based on perceptual fluency (i.e., speed of processing), habituation also causes a perceptual deficit for primed words as they become temporarily disfluent, thereby producing a full crossover to negative priming.

According to the neural habituation model, masking effects arise from nonspecific inhibition between any simultaneously active items, and this masking can occur at all levels of processing. Therefore, even unrelated words produce masking effects, and, furthermore, these masking effects at first increase and then decrease as a function of forward-mask duration as source confusion initially increases but then decreases because of habituation. Because lower levels (e.g., visual features) process information more quickly than do higher levels (e.g., orthographic and lexical–semantic features), it was predicted that the duration producing the most effective forward mask would be longer for unrelated words as compared with pattern masks. Experiments 1 and 3 both produced the predicted U-shaped forward masking effects as a function of mask duration, and both experiments confirmed that pattern masks achieve their greatest effectiveness more rapidly than do unrelated words.

A comparison of Experiments 2 and 3 demonstrated that repetition priming produced both-primed deficits (i.e., trials for which both choice words were equally primed) but associative priming produced both-primed benefits. According the neural habituation model, disfluency due to orthographic habituation is the main factor behind both both-primed deficits as well as the observation of negative priming (Experiment 1) or the elimination of positive priming (Experiment 2) that occurs after long prime viewing durations (Experiment 2 presumably failed to produce negative priming because two more peripheral masked prime words were used, which reduced the magnitude and rapidity of priming). Because semantic priming involves words that do not share orthography, it was expected to produce perceptual both-primed benefits due to top-down support from the common semantics onto dissimilar (nonhabituated) orthography. Competitor models of prime discounting cannot produce these perceptual benefits (Huber et al., 2001; Ratcliff & McKoon, 2001). In keeping with activity-driven habituation as the cause of both-primed deficits, Experiment 2 examined several subliminal priming durations and found that the onset of both-primed deficits with repetition priming lagged behind the onset of priming effects in the target-primed and foil-primed conditions. In contrast, both the ROUSE and multinomial models suppose that both-primed deficits arise from variability in the amount of source confusion, and, therefore, any prime duration sufficient to produce positive (or negative) priming should also produce both-primed deficits. The neural habituation model supposes that habituation is the cause of both-primed deficits, but source confusion (activation) is the cause of both-primed benefits as well as forward-masking effects. In keeping with the latter claim, Experiment 3 observed that the maximum priming benefits occurred for the prime duration that also produced the largest degree of forward masking.

Unlike a traditional neural network, which contains a single differential equation for the dynamic update of activation for each simulated neuron, this model contains two different differential equations for the update of each simulated neuron: The first equation calculates the manner in which activation accrues because of excitatory inputs and inhibition from neighboring cells, and the second equation dictates the manner in which resources deplete (i.e., habituation) and recover. Because each equation is a function of its own past, no direct transformation exists between activation and habituation. In other words, it would be inaccurate to suppose a single underlying construct that automatically moves from activation to habituation. For instance, activation operates on a much faster time scale than habituation and this can be exploited to dissociate these mechanisms. Thus, it should be possible to induce habituation for stimuli that were never fully and explicitly identified. With a near-threshold presentation, the cells that respond to the stimulus will activate only very briefly before activation fades, but even this small amount of activation will induce

a small amount of habituation. Because habituation recovers more slowly as compared with the fading of activation, repeated presentation of near-threshold presentations can accumulate more and more habituation and produce negative priming even for a prime that was never explicitly identified. Recently, this prediction was confirmed in a speeded categorical matching task following multiple subliminal presentations of the category name (Bermeitinger, Frings, & Wentura, 2006; Frings & Wentura, 2005; Wentura & Frings, 2005).

The success of incorporating neural habituation in a cascaded perceptual representation with multiple levels of processing is not limited to the current word-priming experiments. Next, we consider implications of this habituation account for a variety of other paradigms that present items in rapid succession. Habituation is the dominant form of short-term plasticity in many areas of the brain (Thomson & West, 1993; Wang et al., 2006), and so habituation-based discounting should play an important role in many situations.

Priming and the Study of Language

A large literature uses word priming in perceptual identification, speed of naming, or lexical decision tasks to examine the nature of language and semantic representation (e.g., Bavelier, Prasada, & Segui, 1994; Evett & Humphreys, 1981; Ferreira, 2003; Goldinger, 1998; Marcel & Forrin, 1974; McKoon & Ratcliff, 1992; McNamara, 1992, 1994; Meyer & Schvaneveldt, 1971; Meyer et al., 1974; Peressotti & Grainger, 1999; Schlaghecken & Eimer, 2000). However, the current studies, as well as prior reports of similar discounting effects (Hochhaus & Johnston, 1996; Huber, Shiffrin, Lyle, & Quach, 2002; Huber et al., 2001; Huber, Shiffrin, Quach, & Lyle, 2002; Weidemann et al., 2005), place an important cautionary note on all such findings: To compare two or more conditions in terms of priming magnitude, it is critical to keep the strength of discounting constant. Furthermore, it is possible that the strength of discounting is such that it perfectly balances source confusion (i.e., optimal discounting) so as to eliminate priming effects. Therefore, null priming results do not necessarily indicate a lack of relatedness between prime and target. One technique for minimizing discounting is to only use subliminal priming, in which case habituation is expected to play a minor role. However, subliminal priming is tricky for methodological reasons, and a lack of subliminal priming does not necessarily mean that the particular type of priming under consideration is incapable of producing subliminal priming in general. Instead, a lack of subliminal priming might arise from setting prime presentations too far below threshold.

My recommendation is not to fully abandon priming as a measurement tool. Instead, these discounting effects are themselves of interest, and priming can provide a more thorough and reliable understanding of structural representation by taking a dynamic approach to priming that measures priming with both minimal and maximal exposures.

Identification Priming

The phenomenon of repetition blindness can be viewed as a priming paradigm in which prior occurrences of a word make it difficult to individuate subsequent occurrences (Johnston, Hochhaus, & Ruthruff, 2002; Kanwisher, 1987; Kanwisher et al., 1996). However, most priming paradigms involve a type decision (e.g., name the target) rather than a token decision (e.g., name the words in order, listing any repetitions). The forced-choice paradigm includes analogues of both of these paradigms. When given a choice between a repeated target and new foil (the target-primed condition), the decision is akin to the type-based decision common to most priming paradigms where the source of the information does not matter. In this condition, there was a transition from positive to negative priming (Experiment 1). However, when the decision is between a repeated target and a repeated foil (e.g., the both-primed condition of Experiment 2), the question becomes a token question because both choices were seen and the correct answer depends on the source of the presentation. In keeping with the repetition blindness literature, threshold identification in the both-primed condition produces deficits as compared with the neither-primed condition for repetition priming but not associative priming (e.g., Experiment 2, as well as Hochhaus & Johnston, 1996; Huber, Shiffrin, Lyle, & Quach, 2002; Huber et al., 2001; Huber, Shiffrin, Quach, & Lyle, 2002; Weidemann et al., 2005).

Even though the neural habituation model is inherently a type-based model of identification, the dynamics of habituation serve to minimize the effect of prior presentations. Thus, the system is primarily sensitive to the onset of new stimuli, although this occurs at the cost of reduced identification for immediate repetitions. The task for repetition blindness requires report of the entire sequence of words, with this list presumably read out from working memory. This loading into working memory is presumably influenced by the nature of the task. If the task is to keep in mind words in the correct presentation order, including repetitions, then words should only enter working memory when they contain sufficient evidence that they are clearly due to a new presentation rather than lingering activation from an old presentation. Thus, for this task, the loading of working memory requires the same degree of target evidence as is required in the both-primed condition. If, instead, the repetition blindness task was changed to report anything, not worrying about order or repetitions, this would lower the threshold for loading working memory, making it more analogous to the target-primed condition. Indeed, Kanwisher (1987) changed the task in a repetition blindness paradigm to type-based perceptual identification and observed that repeated words were helped rather than hurt.

Moving beyond words to other visual objects, higher level cognitive aftereffects have been demonstrated in the domain of face perception. For instance, a "face space" can be defined in which the opposite from a prime face is specified not in terms of luminance but rather in terms of higher order features such as the width of the face or the distance between the eyes. Within this face space, priming with one face for several seconds shifts the ease of identification away from the prime face and toward the face that lies on the opposite side of the face space (Leopold et al., 2001). In a closely related paradigm, adaptation to a distorted face (e.g., a face in which the eyes have been moved closer together) causes participants to believe that a subsequent nondistorted face is distorted in the opposite direction (Webster & MacLin, 1999). To map out the time course of these face identification effects, my colleague and I used the same paradigm as was used in Experiment 1 but with faces presented at threshold, finding similar transitions from positive to negative face priming (Rieth & Huber, 2005).

Some may consider face habituation and perceived distortions to be low-level perception rather than higher level cognition. More clearly beyond the realm of perception, similar prime-induced

aftereffects are observed by social psychologists studying affect and stereotypes. In this literature, positive priming is referred to as *assimilation* and negative priming is referred to as *contrast* (Martin, 1986; Wegener & Petty, 1997). For instance, in the study of affect and liking responses, Murphy and Zajonc (1993) found that subliminal presentation of affectively valenced pictures (e.g., a smiley face) induce greater liking for a subsequent neutral Chinese ideograph. However, longer, supraliminal exposure to the valenced pictures eliminated this liking effect. Demonstrating that affective valence can fully reverse priming, Glaser and Banaji (1999) examined lexical decisions to affectively valenced target words following presentation of valenced prime words, finding that 150 ms of prime exposure was sufficient to induce faster responding when prime and target were of opposite valence as compared with when they were the same valence (e.g., faster responding to *pleasure* when it followed *torture* versus when it followed *peace*). Note that such results are not easily explained by semantic priming because both congruent and incongruent valence pairs are semantically and/or associatively related. Following up on this study by Glaser and Banaji, my colleagues and I used a speeded liking task with strongly valenced words, finding that the facilitation with brief congruent primes and targets is subsequently eliminated after longer duration primes (Huber, Chun, Parsa, & Winkielman, 2007).

For each of these paradigms, which range from word priming to face priming to affect priming, habituation may provide a parsimonious explanation based on the known dynamics of neural processing. With ongoing exposure to any perceptual or conceptual representation, that representation will eventually habituate, which may serve to reduce or even reverse its effect on an immediate subsequent stimulus.

Recognition Priming

In episodic recognition experiments, immediate repetitions produce a similar transition from positive to negative priming with increasing prime duration. Jacoby and Whitehouse (1989) performed a yes–no recognition experiment in which one condition flashed prime words subliminally prior to participants responding "old" or "new" to a test word, whereas another condition presented primes that were easily seen by participants prior to the test word. A bias effect was observed following subliminal primes such that false alarm rates were higher to unstudied primed test words than unstudied unprimed test words and hit rates were higher to studied primed test words than studied unprimed test words. In contrast, the opposite bias pattern was observed for the easily seen primes such that both false alarms and hits decreased with priming. For subliminal priming, Jacoby and Whitehouse proposed that prime-induced enhanced perceptual fluency was misattributed to episodic familiarity. This account is very similar to the dynamic neural network of Huber and O'Reilly (2003), which likewise supposes lingering activation and enhanced perceptual fluency. However, unlike the neural habituation account, in which discounting is produced automatically, Jacoby and Whitehouse argued that strategic discounting due to prime awareness was responsible for the reversal in the direction of priming following easily seen primes.

Providing evidence against strategic discounting, Joordens and Merikle (1992) replicated both the positive and the negative recognition priming results even though positive priming was achieved with prime durations sufficient to support prime awareness. Nonetheless, using a prime duration that is capable of supporting awareness does not necessarily mean that participants were aware of the primes. Therefore, in another study, my colleagues and I replicated the Joordens and Merikle results but asked participants to explicitly identify which word was primed on a subset of trials, thereby demonstrating that brief primes can produce positive bias effects despite accurate prime identification (Clark, Huber, Curran, & Winkielman, 2007). Furthermore, our studies used forced-choice testing, trial-by-trial feedback, and explicit instructions to discourage strategies. Similar to Experiment 1, we mapped out a gradual nonlinear transition from positive to negative priming for the case of priming recognition memory.

In the Clark et al. (2007) paper, we did not claim that strategic discounting never occurs. Instead, we proposed that neural habituation naturally produces discounting effects and in many situations there is no need to invoke extra strategic processes. This is an argument of parsimony—habituation certainly exists and should be used to explain as much as possible before appealing to more flexible strategic accounts.

Response Priming

Beyond priming the specific identity of an item or word, priming is also seen in terms of responses. For instance, when indicating whether an arrow points to the left or to the right, participants respond more quickly if they answer immediately after a prime arrow indicates the same direction as the target arrow (Eimer, 1995). Several response-priming paradigms involve apparent reductions or even reversals in the direction of priming, although it is difficult to determine whether these are truly response effects or if they are instead the by-product of perceptual habituation. If habituation exists to naturally segment streams of unknown events and reduce source confusion, it could be argued that it is not needed for response behaviors, which are endogenous rather than exogenous (i.e., source confusion is presumably less of an issue with a planned sequence of actions versus an unplanned sequence).

A classic paradigm in response behavior is the Eriksen flanker task (Eriksen & Eriksen, 1974), in which flanking letters that indicate the wrong response slow down responding to a centrally located target letter. Examining the role of priming, Flowers (1980) found that previewing the flankers prior to the onset of the target eliminates flanker effects, but only with sufficiently long preview durations. This was explained by supposing that the response attributes of the previewed flankers were successfully discounted. However, a colleague and I recently performed a series of experiments that identified the locus of these effects in perception rather than in response (Davelaar & Huber, 2006). This was accomplished by including additional preview conditions that did not present any flankers in the final target display. Nearly identical priming effects were seen regardless of whether the flankers were present in the final display. This suggests that preview effects were not due to interactions with the response characteristics of the flankers but, instead, that the preview displays induced a repetition blindness for the flankers.

A similar perception versus response debate exists for the *negative compatibility effect* of Eimer and colleagues (Eimer, 1995; Eimer & Schlaghecken, 2002; Eimer, Schubo, & Schlaghecken, 2002; Schlaghecken & Eimer, 2000). In this paradigm, subliminally presented arrows result in faster responding to target arrows that point in the opposite direction (i.e., negative priming). In

contrast, unmasked prime arrows produce the opposite pattern, revealing facilitation effects with compatible displays (i.e., positive priming). At first blush, these results appear to be exactly opposite the prediction of the neural habituation account based on perception. Indeed, Klapp and Hinkley (2002) proposed a response-based mechanism in which the response inferred by subliminal displays is automatically inhibited. However, the work of Lleras and Enns (2004, 2005, 2006) sheds light on this apparent discrepancy. According to *object substitution masking* (Enns & Di Lollo, 1997), the nature of the masks used to induce subliminal priming can reverse the direction of priming for perceptual reasons. In this account, masks serve to highlight elements that are different than the objects that they mask (i.e., perception is an updating of what is new in each display). Because the masks used in the Eimer studies were composed of arrows pointed in both directions, the masks contrasted against the primes and the perceptual response to the masks was for arrows opposite the primes. Lleras and Enns compared conditions that used these masks, which contained elements of both responses, versus masks that were unbiased in regard to the responses (e.g., a square grid). They replicated the original priming pattern but also found that the pattern of results completely reversed with the unbiased masks, producing compatible effects with subliminal primes.

The substitution masking account of Enns and Di Lollo (1997) is similar to our neural habituation account in some ways, and it may be that neural habituation is the mechanism that underlies object updating (i.e., habituation serves to highlight what's new with each new perceptual object). For instance, consider a brief presentation of a prime arrow (presented for P ms) pointing to the right followed by a mask (presented for M ms) that contains arrows pointing in both directions. In sum, across both the prime and the mask, detection of a rightward arrow receives $P + M$ ms of support and detection of a leftward arrow receives M ms of support. If $P + M$ ms is sufficient to induce habituation, then there will be more habituation for the rightward arrow as compared with the leftward arrow, thus producing negative priming. However, the opposite would be true for a neutral mask because the primed response will only receive P ms of support and the alternative no support, resulting in positive priming. Thus, the use of a mask that contains the same elements as a prime adds additional exposure, which can cause a shift from activation to habituation.

Attentional Paradigms

Beyond the areas of perception, memory, and response, another area of research where habituation may play an important role is attention. In particular, tasks that involve rapid switching of spatial or object-based attention may involve source confusion between attentional representations. For instance, consider a situation in which spatial attention switches between different locations because of exogenous cueing. To reduce the unwanted lingering spatial attention effects of prior locations, it may be useful to deactivate prior locations through habituation to the spatial attention representations associated with those prior locations. Similar to habituation in identification that indicates previous objects that have been sufficiently processed, habituation in spatial attention may indicate previous locations that have been sufficiently processed. This produces spatial attention that readily changes to new locations, but this also causes difficulty in revisiting prior locations (i.e., spatial repetition priming). Indeed, this is exactly what happens with the effect termed *inhibition of return* (Klein, 2000; Posner & Cohen, 1984): Once a location is attended, there is a transient deficit for reallocation of attention to that location.

In addition to spatial attention, object-based attention can exist somewhat independent of spatial position and, more specifically, inhibition of return can apply to objects independent of location (Christ, McCrae, & Abrams, 2002). In the phenomenon of the attentional blink (AB), something similar to habituation of object-based attention lies at the core of the various theoretical accounts of the AB pattern of results. Similar to immediate priming, the AB involves a rapid serial visual presentation (RSVP) stream of items in the same central location at the rate of approximately 10 per second. However, unlike most priming experiments, targets are not defined as the last item presented, but, rather, a target-defining characteristic differentiates between potential targets versus distractors within the RSVP stream. For instance, the original AB paradigm of Raymond, Shapiro, and Arnell (1992) included both letters and numbers in the RSVP stream. The task of the participant was to ignore the numbers and only report the letters once the display was complete. A *blink* is the finding that a second target (T2) is missed for some period of time after accurate identification of a first target (T1).

Considering that there are too many items in the RSVP stream to remember them all, the AB paradigm necessarily involves detection of the target-defining feature to establish when to load items into working memory for later report. Therefore, habituation in this target detection process may be useful for minimizing source confusion that otherwise results in the erroneous loading of subsequent distractors (i.e., there is a need to quickly shut off target detection). But, as with inhibition of return, there is a cost to habituation, and it is difficult to detect a second target for some period of time while the target detector recovers. Considering the large number of AB studies and the wide variety of AB findings (e.g., Nieuwenstein, 2006; Potter, Staub, & O'Connor, 2002; Seiffert & Di Lollo, 1997; Vogel, Woodman, & Luck, 2006), it is not clear whether a model based solely on habituation dynamics can provide a complete account of the AB. Indeed, theoretical accounts of the AB pattern of results appeal to some sort of capacity limitation for the basic deficit as well as some sort of binding mechanism (Bowman & Wyble, 2007) or distractor-initiated closing of attention (Chun & Potter, 1995; Olivers & Watson, 2006). Nevertheless, the proposal of habituation dynamics within target detection is similar to these notions of capacity limitation. Highlighting this similarity, Shapiro, Driver, Ward, and Sorensen (1997) found a close correspondence between immediate priming and the presence or absence of an AB. In these studies, accurate detection of a third target (T3) was examined conditional on detection of the second target (i.e., the "blinked" target). Analogous to Experiment 1, repetition priming between T2 and T3 was negative when T2 was reported (similar to long duration priming), whereas the priming was positive when T2 was not reported (similar to short duration priming).

Another result that is explained by attentional inhibition is the paradigm of *negative priming* (Banks, Roberts, & Ciranni, 1995; Tipper, 1985; Tipper & Driver, 1988), in which a to-be-ignored item on trial n is difficult to respond to as a target on trial $n + 1$ (i.e., slower responses). At the same time, repeating a target from one trial to the next results in facilitation (i.e., faster responses). I use italics to refer to the specific attentional paradigm of *negative priming* rather than any situation where priming is negative. This

pattern of results appears to be the reverse of habituation considering that a minimally processed prime (the ignored distractor) produces negative priming whereas an attended prime produces positive priming. My colleagues and I investigated this situation in a study with a two-prime paradigm similar to Experiment 2, except that participants were given a second task that focused their attention on just one of the two primes (Weidemann et al., 2005). In contrast to the usual results with *negative priming*, we observed positive priming for the ignored prime and negative priming for the attended prime. A complete account of *negative priming* would need to reconcile with these results, as well as the finding that item selection is not a critical aspect of *negative priming* (Milliken, Joordens, Merikle, & Seiffert, 1998). Unlike immediate priming, the *negative priming* paradigm involves a short delay between prime and target because they appear on subsequent trials. This may result in the complete fading of activation, thus leaving only perceptual deficits due to lingering habituation. At the same time, the benefits for responded-to targets in *negative priming* may live within the response system (i.e., it is easier to give the same response on subsequent trials). Such a dynamic interpretation of *negative priming* would be in some ways similar to the episodic retrieval explanation proposed by Neill and colleagues (Neill, 1997; Neill, Valdes, Terry, & Gorfein, 1992), who confirmed that the precise timing of distractors and targets modulate the pattern of results.

Cognitive Neuroscience

One advantage of using a model based on biophysical properties is that it can address neuropsychological, electrophysiological, and hemodynamic results in a straightforward manner. Of these cognitive neuroscience techniques, electrophysiology is best suited for testing the dynamic predictions of the neural habituation model. For the priming paradigm of Experiment 1, a key prediction was that priming effects should exist at the time of the briefly flashed target, which is prior to presentation of the choice words and the initiation of the decision process (see Figure 1). If lingering activation from primes is added to the target response (i.e., source confusion), then the summed response for a cortical region should be greater when the prime and target are different, because different words involve more total representation. This "two different words are more than repeating the same word" prediction arises because imaging techniques cannot measure activation of specific words and instead measure cortical regions—thus, measurements sum across different words if they are concurrently active. Therefore, smaller neural responses with repetitions were expected because of source confusion. With habituation, these *repetition suppression effects* were predicted to diminish because of a reduction in source confusion. In two different event-related electroencephalographic experiments and one event-related magnetoencephalographic experiment, my colleagues and I confirmed the existence of early perceptual repetition effects in response to the target presentation and, furthermore, found the expected changes with prime duration. The neural habituation model with the parameters used in Experiment 1 was applied to these results, providing a simultaneous account of response behavior as well as neural dynamics (Huber, Tian, Curran, O'Reilly, & Woroch, 2007).

In the study of faces and other visual objects, repetition suppression effects are routinely seen with electrophysiology (Campanella et al., 2000; Harris & Nakayama, 2007; Heisz, Watter, & Shedden, 2006; Henson, Rylands, Ross, Vuilleumeir, & Rugg, 2004; Ishai, Bikle, & Ungerleider, 2006; Itier & Taylor, 2004; Jemel, Pisani, Calabria, Crommelinck, & Bruyer, 2003; Kovacs et al., 2006; Martens, Schweinberger, Kiefer, & Burton, 2006), and a few studies have found similar results with visually presented words (Holcomb & Grainger, 2006; Petit, Midgley, Holcomb, & Grainger, 2006). Neural activation and habituation may provide an account of these results, predicting interactions with the nature and timing of prime displays.

Conclusions

Similar to visual aftereffects that produce a positive or negative afterimage as a function exposure duration, cognitive aftereffects exist that can enhance or cause deficits for primed stimuli as a function of prime duration. I propose that lingering neural activity (source confusion) and lingering neural habituation (synaptic depletion) explain the transition from positive to negative priming. Across three priming experiments with perceptual forced-choice word identification, I documented this transition from positive to negative priming and demonstrated that these effects related to forward masking and unbiased deficits (with repetitions) and unbiased benefits (with associative-semantic priming). A dynamic neural network with habituation explained the time course of these priming effects in each experiment. Furthermore, this same model has been applied to a variety of other paradigms, as discussed above. This provides converging evidence that negative priming is a cognitive aftereffect that is beneficial in most situations, serving to reduce source confusion from recent presentations.

References

Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997, January 10). Synaptic depression and cortical gain control. *Science, 275,* 220–224.

Banks, W. P., Roberts, D., & Ciranni, M. (1995). Negative priming in auditory attention. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1354–1361.

Bavelier, D., Prasada, S., & Segui, J. (1994). Repetition blindness between words: Nature of the orthographic and phonological representations involved. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1437–1455.

Bermeitinger, C., Frings, C., & Wentura, D. (2006). Results from repeated masked semantic priming and conclusions to the representation of categorical knowledge [Abstract]. *Journal of Psychophysiology, 20,* 151.

Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review, 114,* 38–70.

Campanella, S., Hanoteau, C., Depy, D., Rossion, B., Bruyer, R., Crommelinck, M., et al. (2000). Right N170 modulation in a face discrimination task: An account for categorical perception of familiar faces. *Psychophysiology, 37,* 796–806.

Castillo, E., Gutierrez, J. M., & Hadi, A. S. (1997). *Expert systems and probabilistic network models.* New York: Springer-Verlag.

Christ, S. E., McCrae, C. S., & Abrams, R. A. (2002). Inhibition of return in static and dynamic displays. *Psychonomic Bulletin & Review, 9,* 80–85.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple-target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 109–127.

Clark, T. F., Huber, D. E., Curran, T., & Winkielman, P. (2007). *Effects of*

repetition priming on recognition memory: Testing a perceptual fluency–disfluency model. Manuscript submitted for publication.

Dasgupta, N., Greenwald, A. G., & Banaji, M. R. (2003). The first ontological challenge to the IAT: Attitude or mere familiarity? *Psychological Inquiry, 14,* 238–243.

Davelaar, E. J., & Huber, D. E. (2006, November). *Preview benefits and deficits in the Eriksen flanker task.* Paper presented at the 47th Annual Meeting of the Psychonomic Society, Houston, Texas.

Duysens, J., Orban, G. A., Cremieux, J., & Maes, H. (1985). Visual cortical correlates of visible persistence. *Vision Research, 25,* 171–178.

Eimer, M. (1995). Stimulus–response compatibility and automatic response activation: Evidence from psychophysiological studies. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 837–854.

Eimer, M., & Schlaghecken, F. (2002). Links between conscious awareness and response inhibition: Evidence from masked priming. *Psychonomic Bulletin & Review, 9,* 514–520.

Eimer, M., Schubo, A., & Schlaghecken, F. (2002). Locus of inhibition in the masked priming of response alternatives. *Journal of Motor Behavior, 34*(1), 3–10.

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

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon identification of a target letter in a nonsearch task. *Perception & Psychophysics, 16,* 143–149.

Evett, L. J., & Humphreys, G. W. (1981). The use of abstract graphemic information in lexical access. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 33*(A), 325–350.

Ferreira, V. S. (2003). The persistence of optional complementizer production: Why saying "that" is not saying "that" at all. *Journal of Memory and Language, 48,* 379–398.

Flowers, J. H. (1980). Priming effects in perceptual classification. *Perception & Psychophysics, 47,* 135–148.

Frings, C., & Wentura, D. (2005). Negative priming with masked distractor-only prime trials: Awareness moderates negative priming. *Experimental Psychology, 52,* 131–139.

Glaser, J., & Banaji, M. R. (1999). When fair is foul and foul is fair: Reverse priming in automatic evaluation. *Journal of Personality and Social Psychology, 77,* 669–687.

Goldinger, S. D. (1998). Signal detection comparisons of phonemic and phonetic priming: The flexible-bias problem. *Perception & Psychophysics, 60,* 952–965.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74,* 1464–1480.

Harris, A., & Nakayama, K. (2007). Rapid face-selective adaptation of an early extrastriate component in MEG. *Cerebral Cortex, 17,* 63–70.

Heisz, J. J., Watter, S., & Shedden, J. M. (2006). Progressive N170 habituation to unattended repeated faces. *Vision Research,* 46, 47–56.

Henson, R. N., Rylands, A., Ross, E., Vuilleumeir, P., & Rugg, M. D. (2004). The effect of repetition lag on electrophysiological and haemodynamic correlates of visual object priming. *Neuroimage, 21,* 1674–1689.

Hochhaus, L., & Johnston, J. C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 355–366.

Hogben, J. H., & Di Lollo, V. (1974). Perceptual integration and perceptual segregation of brief visual stimuli. *Vision Research, 14,* 1059–1069.

Holcomb, P. J., & Grainger, J. (2006). On the time course of visual word recognition: An event-related potential investigation using masked repetition priming. *Journal of Cognitive Neuroscience, 18,* 1631–1643.

Huber, D. E., Chun, W., Parsa, A., & Winkielman, P. (2007). *When is too much of a good thing bad? Testing a Bayesian model of affect priming.* Manuscript in preparation.

Huber, D. E., & Cousineau, D. (2004). A race model of perceptual forced choice reaction time. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 687–692). Hillsdale, NJ: Erlbaum.

Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science: A Multidisciplinary Journal, 27,* 403–430.

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, and 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.

Huber, D. E., Tian, X., Curran, T., O'Reilly, R. C., & Woroch, B. (2007). *The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects.* Manuscript submitted for publication.

Ishai, A., Bikle, P. C., & Ungerleider, L. G. (2006). Temporal dynamics of face repetition suppression. *Brain Research Bulletin, 70,* 289–295.

Itier, R. J., & Taylor, M. J. (2004). Effects of repetition learning on upright, inverted and contrast-reversed face processing using ERPs. *NeuroImage, 21,* 1518–1532.

Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General, 118,* 126–135.

Jemel, B., Pisani, M., Calabria, M., Crommelinck, M., & Bruyer, R. (2003). Is the N170 for faces cognitively penetrable? Evidence from repetition priming of Mooney faces of familiar and unfamiliar persons. *Cognitive Brain Research, 17,* 431–446.

Johnston, J. C., Hochhaus, L., & Ruthruff, E. (2002). Repetition blindness has a perceptual locus: Evidence from online processing of targets in RSVP streams. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 477–489.

Joordens, S., & Becker, S. (1997). The long and short of semantic priming effects in lexical decision. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1083–1105.

Joordens, S., & Merikle, P. M. (1992). False recognition and perception without awareness. *Memory & Cognition, 20,* 151–159.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition, 27,* 117–143.

Kanwisher, N. G., Kim, J. W., & Wickens, T. D. (1996). Signal detection analyses of repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 1249–1260.

Klapp, S. T., & Hinkley, L. B. (2002). The negative compatibility effect: Unconscious inhibition influences reaction time and response selection. *Journal of Experimental Psychology: General, 131,* 255–269.

Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Sciences, 4,* 138–147.

Kotulak, J. C., & Schor, C. M. (1986). The accommodative response to subthreshold blur and to perceptual fading during the Troxler phenomenon. *Perception, 15,* 7–15.

Kovacs, G., Zimmer, M., Banko, E., Harza, I., Antal, A., & Vidnyanszky, Z. (2006). Electrophysiological correlates of visual adaptation to faces and body parts in humans. *Cerebral Cortex, 16,* 742–753.

Kučera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English.* Providence, RI: Brown University Press.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience, 4,* 89–94.

Lleras, A., & Enns, J. T. (2004). Negative compatibility or object updating? A cautionary tale of mask-dependent priming. *Journal of Experimental Psychology: General, 133,* 475–493.

Lleras, A., & Enns, J. T. (2005). Updating a cautionary tale of masked priming: Reply to Klapp (2005). *Journal of Experimental Psychology: General, 134,* 436–440.

Lleras, A., & Enns, J. T. (2006). How much like a target can a mask be? Geometric, spatial, and temporal similarity in priming: A reply to Schlaghecken and Eimer (2006). *Journal of Experimental Psychology: General, 135,* 495–500.

Long, G. M., Toppino, T. C., & Mondin, G. W. (1992). Prime time: Fatigue and set effects in the perception of reversible figures. *Perception & Psychophysics, 52,* 609–616.

Lukatela, G., Eaton, T., Lee, C. H., Carello, C., & Turvey, M. T. (2002). Equal homophonic priming with words and pseudohomophones. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 3–21.

Lukatela, G., Frost, S. J., & Turvey, M. T. (1998). Phonological priming by masked nonword primes in the lexical decision task. *Journal of Memory and Language, 39,* 666–683.

Marcel, T., & Forrin, B. (1974). Naming latency and repetition of stimulus categories. *Journal of Experimental Psychology, 103,* 450–460.

Martens, U., Schweinberger, S. R., Kiefer, M., & Burton, A. M. (2006). Masked and unmasked electrophysiological repetition effects of famous faces. *Brain Research, 1109,* 146–157.

Martin, L. L. (1986). Set/reset: Use and disuse of concepts in impression formation. *Journal of Personality and Social Psychology, 51,* 493–504.

Masson, M. E. J., & Borowsky, R. (1998). More than meets the eye: Context effects in word identification. *Memory & Cognition, 26,* 1245–1269.

McKoon, G., & Ratcliff, R. (1992). Spreading activation versus compound cue accounts of priming: Mediated priming revisited. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1155–1172.

McNamara, T. P. (1992). Theories of priming: I. Associative distance and lag. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1173–1190.

McNamara, T. P. (1994). Theories of priming: II. Types of primes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 507–520.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90,* 227–234.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1974). Functions of graphemic and phonemic codes in visual word recognition. *Memory & Cognition, 2,* 309–321.

Milliken, B., Joordens, S., Merikle, P. M., & Seiffert, A. E. (1998). Selective attention: A reevaluation of the implications of negative priming. *Psychological Review, 105,* 203–229.

Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology, 64,* 723–739.

Neill, W. T. (1997). Episodic retrieval in negative priming and repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1291–1305.

Neill, W. T., Valdes, L. A., Terry, K. M., & Gorfein, D. S. (1992). Persistence of negative priming: II. Evidence for episodic trace retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 993–1000.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1994). *The University of South Florida word association, rhyme and word fragment norms.* Unpublished manuscript.

Nelson, S. B., Varela, J. A., Sen, K., & Abbott, L. F. (1997). Functional significance of synaptic depression between cortical neurons. In J. Bowers (Ed.), *Computational neuroscience* (pp. 429–434). New York: Plenum Press.

Nieuwenstein, M. R. (2006). Top-down controlled, delayed selection in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 973–985.

Olivers, C. N. L., & Watson, D. G. (2006). Input control processes in rapid serial visual presentations: Target selection and distractor inhibition. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 1083–1092.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain.* Cambridge, MA: MIT Press.

Peressotti, F., & Grainger, J. (1999). The role of letter identity and letter position in orthographic priming. *Perception & Psychophysics, 61,* 691–706.

Petit, J. P., Midgley, K. J., Holcomb, P. J., & Grainger, J. (2006). On the time course of letter perception: A masked priming ERP investigation. *Psychonomic Bulletin & Review, 13,* 674–681.

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. Bouwhuis (Eds.), *Attention and performance X: Control of language processes* (pp. 531–556). Hillsdale, NJ: Erlbaum.

Potter, M. C., Staub, A., & O'Connor, D. H. (2002). The time course of competition for attention: Attention is initially labile. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 1149–1162.

Ratcliff, R., Hockley, W., & McKoon, G. (1985). Components of activation: Repetition and priming effects in lexical decision and recognition. *Journal of Experimental Psychology: General, 114,* 435–450.

Ratcliff, R., & McKoon, G. (1994). Retrieving information from memory: Spreading-activation theories versus compound-cue theories. *Psychological Review, 101,* 177–184.

Ratcliff, R., & McKoon, G. (2001). A multinomial model for short-term priming in word identification. *Psychological Review, 108,* 835–846.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18,* 849–860.

Rieth, C. A., & Huber, D. E. (2005). Using a neural network model with synaptic depression to assess the dynamics of feature-based versus configural processing in face identification. In B. G. Bara, L. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 1856–1861). Hillsdale, NJ: Erlbaum.

Schlaghecken, F., & Eimer, M. (2000). A central–peripheral asymmetry in masked priming. *Perception & Psychophysics, 62,* 1367–1382.

Schooler, L. J., & Anderson, J. R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology, 32,* 219–250.

Seiffert, A. E., & Di Lollo, V. (1997). Low-level masking in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 1061–1073.

Shapiro, K., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science, 8,* 95–100.

Shively, F. (1973). New afterimage (color contrast afterimage). *Perception & Psychophysics, 13,* 525–526.

Smith, L., & Klein, R. (1990). Evidence for semantic satiation: Repeating a category slows subsequent semantic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 852–861.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74*(11, Whole No. 498).

Stapel, D. A., & Winkielman, P. (1998). Assimilation and contrast as a function of context–target similarity, distinctness, and dimensional relevance. *Personality and Social Psychology Bulletin, 24,* 634–646.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience, 19,* 109–139.

Taylor, G. A., & Chabot, R. J. (1978). Differential backward masking of

words and letters by masks of varying orthographic structure. *Memory & Cognition, 6,* 629–635.

Thomson, A. M., & West, D. C. (1993). Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. *Neuroscience, 54,* 329–346.

Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 37*(A), 571–590.

Tipper, S. P., & Driver, J. (1988). Negative priming between pictures and words in a selective attention task: Evidence for semantic processing of ignored stimuli. *Memory & Cognition, 16,* 64–70.

Tsodyks, M. V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences, USA, 94,* 719–723.

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 1436–1451.

Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J. Y., & Goldman-Rakic, P. S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nature Neuroscience, 9,* 534–542.

Webster, M. A., & MacLin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review, 6,* 647–653.

Wegener, D. T., & Petty, R. E. (1997). The flexible correction model: The role of naive theories of bias in bias correction. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 29, pp. 141–208). San Diego, CA: Academic Press.

Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2005). Spatiotemporal confusion and compensation in visual word perception. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 40–61.

Wentura, D., & Frings, C. (2005). Repeated masked category primes interfere with related exemplars: New evidence for negative semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 108–120.

Yee, P. L., Santoro, K. E., Grey, A. L., & Woog, V. (2000). Deeper processing at target selection increases the magnitude of negative priming. *Memory & Cognition, 28,* 1429–1436.

Appendix A

Forced Choice Accuracy in Experiment 1

Prime duration (ms)	Pattern mask		Unrelated word		Target primed		Foil primed	
	M	*SE*	*M*	*SE*	*M*	*SE*	*M*	*SE*
17	.879	.017	.807	.021	.921	.015	.616	.030
50	.799	.024	.707	.018	.875	.020	.393	.025
150	.807	.021	.674	.021	.776	.023	.488	.029
400	.820	.024	.809	.019	.594	.031	.809	.023
2,000	.845	.020	.809	.019	.550	.028	.832	.020

Appendix B

Forced Choice Accuracy in Experiment 2

Prime duration (ms)	Unrelated words		Both primed		Target primed		Foil primed	
	M	*SE*	*M*	*SE*	*M*	*SE*	*M*	*SE*
14	.896	.026	.883	.027	.892	.026	.913	.021
43	.896	.023	.888	.023	.922	.018	.850	.031
129	.759	.029	.674	.029	.805	.026	.614	.037
386	.701	.032	.627	.032	.744	.030	.559	.048

Appendix C

Forced Choice Accuracy in Experiment 3

Prime duration (ms)	Pattern mask		Unrelated word		Both primed	
	M	*SE*	*M*	*SE*	*M*	*SE*
57[a]	.873	.013	.668	.016	.716	.015
29[a]	.874	.014	.811	.013	.821	.013
0[b]	.863	.012	.863	.012	.863	.012
29	.777	.015	.833	.013	.858	.014
57	.682	.015	.741	.016	.773	.014
114	.702	.015	.694	.016	.745	.017
457	.711	.015	.695	.017	.736	.018
1,829	.757	.016	.753	.015	.757	.016

[a] These conditions presented the prime after the target for the indicated duration.
[b] No prime (0 ms) is a single condition, but it is shown three times for comparison with each of the three priming conditions.

Received January 11, 2007
Revision received September 25, 2007
Accepted September 28, 2007 ■

ORDER FORM

Start my 2008 subscription to the ***Journal of Experimental Psychology: General*** ISSN: 0096-3445

___ $52.00, **APA Member/Affiliate** _____
___ $84.00, **Individual Nonmember** _____
___ $300.00, **Institution** _____
In DC add 5.75% / In MD add 6% sales tax _____
Total Amount Enclosed $ _____

Subscription orders must be prepaid. (Subscriptions are on a calendar year basis only.) Allow 4-6 weeks for delivery of the first issue. Call for international subscription rates.

American Psychological Association

SEND THIS ORDER FORM TO:
American Psychological Association
Subscriptions
750 First Street, NE
Washington, DC 20002-4242

Or call **800-374-2721**, fax **202-336-5568**.
TDD/TTY **202-336-6123**.
For subscription information, e-mail:
subscriptions@apa.org

◯ **Check enclosed** (make payable to APA)

Charge my: ◯ VISA ◯ MasterCard ◯ American Express

Cardholder Name _____
Card No. _____ Exp. Date _____

_____
Signature (Required for Charge)

BILLING ADDRESS:
Street _____
City _____ State _____ Zip _____
Daytime Phone _____
E-mail _____

MAIL TO:
Name _____
Address _____
_____
City _____ State _____ Zip _____
APA Member # _____ *XGEA08*