The Dynamics of Integration and Separation: ERP, MEG, and Neural Network Studies of Immediate Repetition Effects

David E. Huber
University of California, San Diego

Xing Tian
University of Maryland

Tim Curran, Randall C. O'Reilly, and Brion Woroch
University of Colorado at Boulder

This article presents data and theory concerning the fundamental question of how the brain achieves a balance between integrating and separating perceptual information over time. This theory was tested in the domain of word reading by examining brain responses to briefly presented words that were either new or immediate repetitions. Critically, the prime that immediately preceded the target was presented either for 150 ms or 2,000 ms, thus examining a situation of perceptual integration versus one of perceptual separation. Electrophysiological responses during the first 200 ms following presentation of the target word were assessed using electroencephalography (EEG) and magnetoencephalography (MEG) recordings. As predicted by a dynamic neural network model with habituation, repeated words produced less of a perceptual response, and this effect diminished with increased prime duration. Using dynamics that best accounted for the behavioral transition from positive to negative priming with increasing prime duration, the model correctly predicted the time course of the event-related potential (ERP) repetition effects under the assumption that letter processing is the source of observed P100 repetition effects and word processing is the source of observed N170 repetition effects.

Keywords: repetition priming, ERP, MEG, habituation, neural network

The brain must achieve a balance between integrating and separating information over time. Over short durations (e.g., tens of milliseconds), the brain needs to integrate sensory information to provide a stable and interpretable percept; but, over longer durations, it needs to distinguish between distinct sensory events. Behaviorally, these phenomena have been studied using short-term priming experiments, where two stimuli are presented close in time and the impact of the first (prime) stimulus on perception of the second (target) stimulus is measured (e.g., Evett & Humphreys, 1981; Meyer & Schvaneveldt, 1971).

Using identical prime-target stimuli (i.e., repetition priming) and a forced-choice test of words flashed at perceptual threshold, recent studies have observed enhanced identification of repeated targets following short prime durations (e.g., less than 200 ms), but impaired identification following longer prime durations (Huber, 2008; Huber, Shiffrin, Lyle, & Quach, 2002; Huber, Shiffrin, Lyle, & Ruys, 2001; Huber, Shiffrin, Quach, & Lyle, 2002; Weidemann, Huber, & Shiffrin, 2005, 2008). Such findings help clarify the large and varied literature regarding short-term priming facilitations and deficits (e.g., Eimer, 1995; Hochhaus & Johnston, 1996; Humphreys, Besner, & Quinlan, 1988; Marcel, 1983; Masson & Borowsky, 1998; Meyer & Schvaneveldt, 1971; Meyer, Schvaneveldt, & Ruddy, 1974), suggesting that the extent of prime processing is a critical factor that may potentially eliminate or even reverse priming effects.

Such reversals with prolonged presentations or with increased prime salience are found with many other paradigms. For instance, interpretation of a visual stimulus that is ambiguous in the depth plane is biased towards a primed interpretation following an unambiguous prime, but this finding reverses to a bias against the primed interpretation when the unambiguous prime is seen for many seconds (Long, Toppino, & Mondin, 1992). A similar result is found with face perception. Leopold, O'Toole, Vetter, and Blanz (2001) created a "face space" using morph technology, finding that prolonged (5 s) presentation of an antiface enhanced perception of a target face (see Rhodes et al., 2005, for a recent review of face adaptation effects). The inability to adequately perceive a repeated stimulus is also found in the phenomenon termed "repetition blindness," in which the second occurrence of word within a sentence is not noticed (Kanwisher, 1987; Kanwisher, Kim, & Wickens, 1996; Whittlesea, Dorken, & Podrouzek, 1995). Similar

reversals after extended processing are also found in less perceptual domains, such as attentional processing for inhibition of return (Klein, 2000; Posner & Cohen, 1984; Samuel & Kat, 2003) and reversals in evaluative priming (Glaser & Banaji, 1999; Murphy & Zajonc, 1993). Huber and O'Reilly (2003) reviewed these paradigms and others, relating them to forced-choice perceptual identification, as well as to the neural dynamics that may underlie all these phenomena. Subsequently, Huber (2008) tested this neural account behaviorally with a range of prime durations while examining repetition priming, associative priming, and masking effects. The current studies take this same account and test neural behavior as predicted from response behavior.

In the reported studies, we used forced-choice perceptual identification for words to investigate these dynamics, but the theory we tested has implications for many tasks that involve the immediate effect of one stimulus on a subsequent stimulus (see Huber, 2008, for a review of these implications). The current paradigm modifies the original perceptual identification task of Humphreys et al. (1988) by using forced-choice testing to reveal the costs as well as the benefits of priming through conditions that prime the wrong alternative versus conditions that prime the correct alternative. As seen in Figure 1, the participants' task was to identify the briefly flashed target word (e.g., PATCH) presented in the center of the screen immediately following presentation of one (Experiment 1) or two (Experiments 2 and 3) prime words. Target flash

durations were set individually at the perceptual threshold for each participant such that accuracy was approximately 75% in forced-choice identification (e.g., a choice between HURRY and PATCH). In the *repeated* condition, the target was a repetition of a prime word (also known as the target-primed condition), and in the *novel* condition, the target was a word previously unseen in the experiment. A novel target flash could be further broken down into two additional behavioral conditions, depending on whether neither choice word was primed (a true baseline condition, also known as the neither-primed condition, which is not shown in Figure 1) or whether the incorrect foil choice word was primed (a condition that assesses the costs of priming, also known as the foil-primed condition, which is labeled *novel* in Figure 1). The current experiments only included this foil-primed case that assesses the costs of priming; but, for reasons of consistency, we nonetheless refer to this as the novel condition, considering that we examined electrophysiological responses to the briefly flashed target word prior to the onset of the choices (thus, nothing was primed at this point within the sequence of events). Nevertheless, it is important to keep in mind that the novel condition still involved priming, although it took the form of priming a choice alternative (i.e., decisional priming) rather than priming the target.

Priming may result in an interactive effect, for instance if more information is extracted more quickly from the briefly flashed target. A positive interaction between prime and target is defined
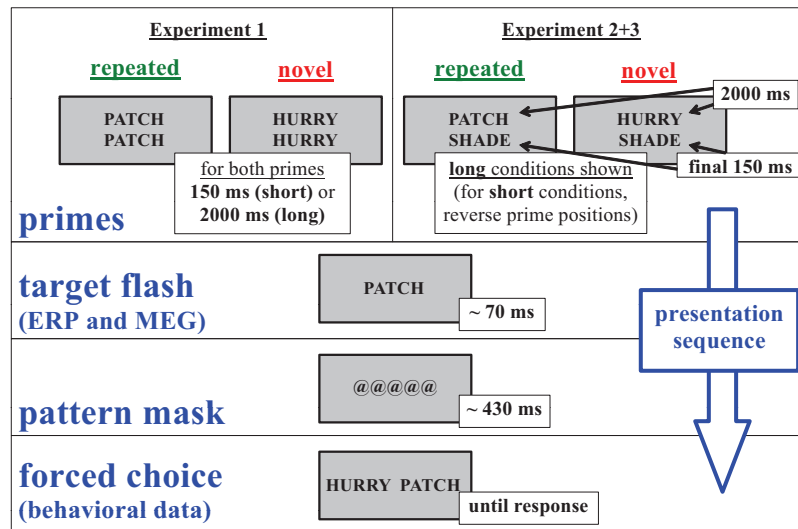

Figure 1. Presentation sequence for the reported experiments. Behavioral responses are collected at the time of the forced-choice presentation (decision) but the reported experiments examined electrophysiological responses at the time of the target flash (perception). Participants attempted to identify which choice word was the same as the briefly flashed target word. Target flash duration was set individually to place performance at 75%, which resulted in a duration of 70 ms on average. The pattern mask duration was set so the delay between target onset and forced choice onset was 500 ms. All experiments compare conditions where the target repeats a prime (repeated) versus conditions where the target is new (novel) but the incorrect foil choice word repeats a prime (i.e., a foil-primed condition that assesses the costs of priming). Experiment 1 used a single prime word presented for 150 ms (short) or 2,000 ms (long), with the same prime word appearing simultaneously both above and below fixation. Experiments 2 and 3 used two different prime words, with the upper prime word presented for 2,000 ms and the lower presented for the final 150 ms prior to the target flash. One of the two choice words was a repetition of either the upper prime (long) or the lower prime (short). Participants were instructed that there could be no effective strategy in relation to the prime words considering that the correct target word was just as often a repetition of a prime as the incorrect foil word. Trial-by-trial feedback reinforced this claim.

as larger performance gains when priming the target compared to performance losses when priming the foil. Alternatively, priming may result in an additive effect, changing performance without changing the information gained from the target flash, which results in equal-sized performance gains and losses when priming the target versus priming the foil, both compared to the neither-primed condition in which neither choice is primed. If the main effect of priming is one of spatial/temporal source confusion, with the primes mistaken for the briefly flashed target, this will be revealed as an additive effect, with increased choice accuracy in the repetition condition but decreased choice accuracy in the novel condition that primes the foil. For example, if both the prime and target are the word PATCH, then this might result in a higher probability of correctly choosing PATCH. In contrast, if the prime is HURRY and the target is PATCH, this might result in an increasing probability of incorrectly choosing HURRY (i.e., lower performance compared to no priming).

The reported experiments do not include a truly unprimed condition (i.e., prime words unrelated to both choice words) against which these two priming conditions are compared, but previous research (Hochhaus & Johnston, 1996; Huber, 2008; Huber, Shiffrin, Lyle, et al., 2002; Huber et al., 2001; Huber, Shiffrin, Quach, et al., 2002; Masson & Borowsky, 1998; Weidemann et al., 2005, 2008), demonstrated that priming is largely additive, with approximately symmetric costs and benefits (although see Huber, 2008, for conditions that demonstrated small interactive deficits with repetition priming and small interactive benefits with associative-semantic priming). It is tempting to consider a decisional interpretation of such additive priming effects in which participants simply wait for the choice words and choose the word that is recognized as the prime. If applied to those trials where the target was not directly perceived, this would help performance in the target-primed (repeated) condition but hurt performance in the foil-primed (novel) condition, both compared to the baseline neither-primed condition. The primary goal of this work is to differentiate between such a decisional account versus a perceptual account by examining brain responses prior to the presentation of the choice words that define the choice alternatives.

Thus, electrophysiology is ideally suited to covertly measure when in time the priming effect occurs: At the time of the choice words (decision) or at the time of the target flash (perception).

With prime durations ranging from subliminal to several seconds (see Figure 2), priming remains largely additive in nature (Huber, 2008), although the pattern of prime induced costs and benefits changes direction from positive to negative priming (Huber, 2008; Huber, Shiffrin, Lyle, et al., 2002; Huber et al., 2001; Huber, Shiffrin, Quach, et al., 2002; Weidemann et al., 2005, 2008), and this is taken as the signature pattern of the transition from integration to separation of perceptual events (Huber & O'Reilly, 2003). As seen in Figure 2, these effects are very large, and 50 ms of prime viewing can result in reliably choosing the wrong answer more often than the correct answer for the novel condition that primes the incorrect choice. In other words, people are confused regarding what was presented when and where and often fully mistake the prime for the target. Nevertheless, after a few hundred milliseconds of prime viewing, this pattern completely reveres and now the tendency is *not* to choose whatever has been primed.

These priming reversals are explained as the discounting of the prime in order to reduce source confusion. Furthermore, there is ample evidence that this discounting is implicit in nature rather than due to strategic bias. For instance, these effects exist across manipulations of the proportion of primed trials (Weidemann et al., 2008). Additionally, discounting is ineffective in some conditions but not others, even though the primes are equally salient and even though these different conditions appear in a mixed design (see Huber, Shiffrin, Lyle, et al., 2002, for an explanation of these conditions). Demonstrating these effects, prime durations that would otherwise produce negative priming, instead produced positive priming when the decision is between very similar choice words (Huber et al., 2001), when no target is presented on some trials (Huber, Shiffrin, Lyle, et al., 2002; Weidemann et al., 2008), and when the similarity between primes and primed choice words is very low (Huber, Shiffrin, Lyle, et al., 2002; Huber et al., 2001). Nevertheless, concluding that these effects are implicit does not imply that discounting is perceptual in its origin. Therefore, the
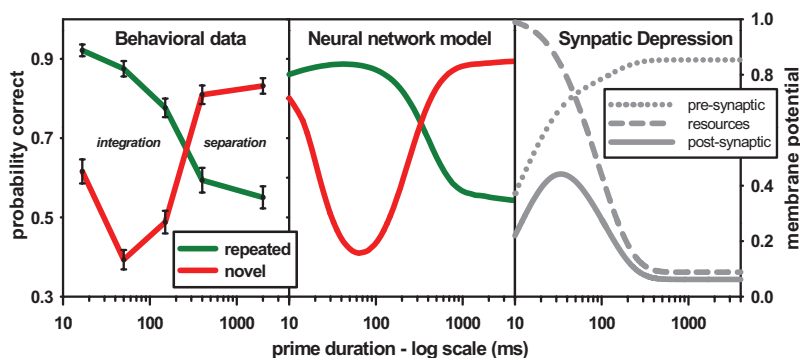
Figure 2. Behavioral data (Huber, 2008) and neural network behavior (Huber & O'Reilly, 2003) using the same paradigm as Experiment 1 but with the inclusion of additional prime durations. The right-hand panel exemplifies the role of neural habituation through synaptic depression, resulting in postsynaptic potentials that reach a maximum and subsequently reduce to a low asymptote following extended prime presentations. This deactivation due to loss of synaptic resources produces a transition from integration to separation between prime and target flash.

current experiments examined early electrophysiological responses to the briefly flashed target word, which occurs prior to presentation of the choice options (i.e., prior to the decision process), to measure the perceptual basis of variations in priming as a function of prime duration. If these perceptual neural responses modulate in the manner predicted by behavioral priming, this will provide converging evidence that discounting can directly change perceptual processes.

A Perceptual Account of Priming Based on Neural Habituation

According to the neural habituation model of Huber and O'Reilly (2003), which places these priming effects in the perceptual response, source confusion (integration) between prime and target explains the initial tendency to mistakenly blend prime and target, thus resulting in a greater preference (i.e., additive effect) for primed words. However, with excessive prime processing, discounting (separation) results in small or missing perceptual responses to primed words and, therefore, an apparent preference against primed words. Huber and O'Reilly hypothesized that the electrophysiological properties of neurons that integrate excitatory input over time, combined with transient neural habituation, explain this transition from integration to separation. According to their account, these effects naturally and automatically result from the neural dynamics of perception. Simulating a dynamic neural network with these properties, their theory explained a wide range of relevant behavioral data.

A key assumption underlying this account is that participants monitor the perceptual response to each choice word in deciding which option was just presented as the target. In other words, the perceptual fluency of each choice word is used to guide responding. This is a sensible response strategy considering that the target is presented at the perceptual threshold and is not often explicitly identified. Source confusion from the primes boosts the perceptual response to a primed choice word, whereas discounting reduces the perceptual response to a primed choice word. This reduction in source confusion not only minimizes the perceptual influence of the prime on the target, but it additionally results in a perceptual disfluency such that the perceptual response to a primed word occurs more slowly compared to an unprimed word. Thus, because performance is based on fluency, there is a real deficit in processing a primed choice word, and this can result in a full reversal to the direction of behavioral priming. In keeping with this claim, reaction times are found to be faster when choosing a primed word in the positive priming situation but are found to be slower when choosing a primed word in the negative priming situation (Huber & Cousineau, 2004).

The transition from integration to separation is a useful mechanism for clearing prior activation, allowing unobstructed perceptual processing of new input. It greatly reduces source confusion and eliminates the tendency to blend together subsequent perceptual events. As seen in Figure 2, priming the incorrect choice (i.e., the novel condition) is actually beneficial, provided that the prime is viewed sufficiently long prior to the briefly flashed target. Once the perceptual system has fully identified and dealt with a perceptual object, that object is ruled out as the potential cause of subsequent perceptions (temporary depression), thereby making it easier to perceive different subsequent objects (i.e., novelty detec-

tion). However, such a clearing of activation comes at a cost for those occasions when a closely related or identical item subsequently appears (i.e., the repeated condition in Figure 2). Assuming that subsequent items are more likely to be new rather than immediate repetitions, this can be viewed as an adaptive mechanism for reducing the unwanted effects of source confusion between perceptual events that are close in time. Besides reducing source confusion, this separation serves to highlight novelty, causing the new aspects of perceptual input to pop out (i.e., temporal contrast), which is also adaptive in terms of drawing attention to new, possibly threatening, events.

This bottom-up account of repetition priming deficits assumes that separation occurs automatically during perception, affecting processing of the target prior to the choice process. In contrast, top-down accounts appeal to discounting in the decision process (e.g., Jacoby & Whitehouse, 1989) or implicit decisional processes based on optimality. For instance, the responding optimally with unknown sources of evidence (ROUSE) Bayesian decision model (Huber et al., 2001) assumes that perceptual processing is unchanged and that the detrimental effect of an extended prime presentation is a decision effect in response to the choice words. Thus, the account based on perceptual separation predicts effects at the time of the briefly presented target, whereas accounts based on decision discounting place the effect at the time of the choice words.

Behavioral examination of target processing requires some sort of decision process or additional presentation, and, therefore, it is difficult to differentiate between these accounts based upon overt response behavior. In contrast, neuroimaging provides a passive method for covertly monitoring the response to the target word without eliciting a decision. In the three reported experiments, we differentiated between a perceptual account versus a decisional account by examining perceptual (i.e., <200 ms) brain responses to the briefly flashed target word, which is presented prior to the choice words. The key qualitative prediction, as outlined below, is for less of a perceptual response for repeated targets, but this repetition deficit should lessen in magnitude with increased prime duration. Next, we describe the representational and dynamic assumptions of the neural habituation model that we are testing. We explain the manner in which it handles behavioral data and the manner in which it was used to produce a priori predictions for neural repetition effects.

Model Architecture and Dynamics

Full specification of the neural network model was previously reported by Huber and O'Reilly (2003) and a summary is provided here. As seen in Figure 3A, the bottom level of the model encodes visual input at different visual locations (e.g., visual lines), whereas the middle and top levels encode progressively abstract higher level representations, such as with identification of letters and words regardless of screen location. Therefore, these higher levels suffer both from spatial source confusion and from temporal source confusion, although habituation serves to reduce temporal source confusion. In a series of experiments that parametrically varied prime duration, masking, and the type of priming (repetition versus associative), Huber (2008) tested not only the dynamics of this model but also this particular architecture of lines, letters, and word-level representations, with top-down feedback from the word
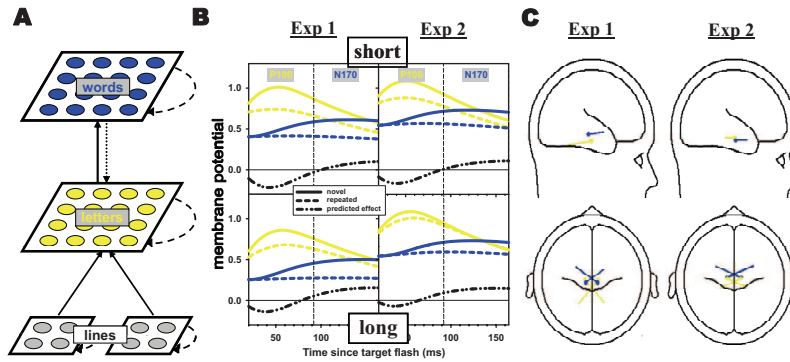
Figure 3. A: The three-level model of Huber and O'Reilly (2003). B: Event-related potential (ERP) repetition effect predictions following one prime, Experiment 1 (Exp 1), or two primes, Experiment 2 (Exp 2), in the short and long conditions using previously published parameters (same parameters as Figure 2). C: Equivalent source dipole positions and orientations that best accounted for the observed ERP results of Experiments 1 and 2. Level 1 is sensitive to lines at specific visual locations. Level 2 responds to the same sequence of letters regardless of where those letters appear. Level 3 responds to known words regardless of where they appear. The yellow and blue lines in Panel B present membrane potential summed across all units of the letter (yellow) or word (blue) levels. ERP repetition effect predictions are derived by assuming that the letter level produces the P100 peak and that the word level produces the N170 peak. Thus, the black line at the bottom of each graph is the difference of the blue lines subtracted from the difference of the yellow lines, which corresponds to the assumption of equal and oppositely oriented dipoles for the letter and word levels. There is a smaller voltage for repetitions during the time period of the P100 (smaller P100s) and a larger voltage for repetitions during the time period of the N170 (smaller N170s). Furthermore, the P100 effect is predicted to diminish and the N170 effect is predicted to increase following long-duration primes. This N170 effect as a function of prime duration is due to the lingering P100 letter-level effect overwriting the N170 word-level effect following a short-duration prime. Only following a long-duration prime is the letter-level effect sufficiently diminished to fully reveal the underlying word level N170 effect. This is particularly true for Experiment 2 because two prime words produce greater inhibition at the letter level. The results of Panel C confirmed the assumption of equal magnitudes for each level (the length of the yellow and blue lines) and opposite orientations for each level (the direction of the yellow and blue lines).

level back to the letter level. This architecture is largely inspired by the original interactive-activation model of word reading of McClelland and Rumelhart (1981), which explains the word superiority effect, except for the inclusion of habituation.

In simulations with this model, words are presented in the same temporal and spatial manner as in the experiments and activation is updated every millisecond. Each simulated neuron includes excitatory, inhibitory, and leak currents, as dictated by the electrical and diffusion forces operating in real neurons. The effect of inhibitory interneurons is simulated through local inhibition between all units within a patch of units (i.e., within the boxes of Figure 3A). The critical activation variable of these neurons is membrane potential rather than actual spiking behavior, and each simulated neuron can be thought of as representing a large number of spiking neurons that have similar inputs and outputs, with membrane potential proportional to the average firing rate of the assembly.

When a neuron produces an action potential and subsequent neurotransmitter release, the effect of the released neurotransmitter on the receiving (i.e., postsynaptic) cell lasts for a period of time ranging from tens to hundreds of milliseconds. Due to these lingering effects, the excitation from prior action potentials adds with subsequent action potentials, yielding even greater postsynaptic depolarization (e.g., the voltage potential across the membrane of the receiving cell), and, therefore, greater probability of

producing action potentials in the receiving cell. As a stimulus is presented longer, or repeated in rapid succession, more total activation accrues due to this integration process. The model captures this integration with a Hodgkin-Huxley differential equation for the gradual increase in postsynaptic membrane potential in the face of ongoing excitatory input (O'Reilly & Munakata, 2000). This equation includes a leak current, such that membrane potential gradually returns to baseline in the absence of excitatory input. Because the leak current is relatively weak, the return to baseline takes some time, resulting in lingering persistent activation past removal of input. This enables the integration of excitatory input over short time periods.

There are other dynamic properties of synaptic transmission operating on even slower time scales that serve to eventually counteract the accumulation of excitatory responses. In other words, neural responses tend to habituate with ongoing excitation due to a variety of mechanisms. Because these processes are slower, postsynaptic neural activity achieves a maximum and then falls to a low asymptotic level in the presence of prolonged exposure to a stimulus. These counteracting processes are collectively referred to as synaptic depression, which arises from a variety of cellular mechanisms that diminish the postsynaptic effect of each additional presynaptic (i.e., sending) action potential. These cellular mechanisms include both the depletion of the resources necessary for effective synaptic signaling as well as

more direct forms of inhibition. This depletion is temporary in nature and is driven by recent activity. For instance, calcium is needed for the release of the vesicles that contain neurotransmitter, and over time, calcium is gradually depleted as the presynaptic cell continues to release neurotransmitter. Other counteracting processes work on the postsynaptic side, serving to desensitize the postsynaptic cell to excitatory input. Still other processes work through slowly accruing inhibitory currents. Regardless of their source, these counteracting processes have been collectively termed synaptic depression and quantified through electrophysiological experiments and mathematical models (e.g., Abbott, Varela, Sen, & Nelson, 1997; Tsodyks & Markram, 1997). Huber and O'Reilly (2003) used the equation of Abbott et al. (1997) to dynamically modify the strength of excitatory input between sending and receiving cells. It is important to note that synaptic depression is transient and is not the same as the long-term connection strength (i.e., the weight) in a traditional neural network. Instead, synaptic depression is a dynamically varying term that specifies the transient efficacy of the connection. It is also important to note that synaptic depression does not have a specific time course because it is the ongoing use of synaptic resources that drives depletion. If there is less presynaptic activity, then resources will be depleted at a slower rate, and synaptic depression will take place more slowly.

The dynamics of habituation through synaptic depression are exemplified in the right-hand panel of Figure 2, which shows the separate variables of presynaptic membrane potential (i.e., activation), synaptic resources, and the resultant postsynaptic depolarization (i.e., output). Postsynaptic depolarization is equal to the multiplication of membrane potential and synaptic resources, which can be thought of as the probability of spiking multiplied by the resultant depolarization associated with each spike. Although the same parameters are used for every simulated neuron across all three levels of the network, the integration time constant is set to smaller values for higher levels of the network such that the letter and word representations are slower to activate and slower to habituate. This is sensible considering the rapidly changing nature of visual input as contrasted with the need to slowly extract and maintain meaning across words and sentences.

Modeling Behavioral Data

The right-hand panel of Figure 2 shows a Level 1 neuron (e.g., response to visually presented lines) with parameters that produced the accuracy results seen in the middle panel. Because the postsynaptic depolarization shown in the right-hand panel is the input to Level 2, this activity dependent "n-shaped" input to letter processing helps explain the seemingly complex behavioral pattern, depending on whether the prime is the same as the target ("n-shaped" accuracy) or the same as the foil ("u-shaped" accuracy). Because the target flash duration is set at the identification threshold in this paradigm, partial activation from the target is assessed in the decision process according to which choice word activates first. Residual activation from the target flash supports accurate responding but residual activation from the prime is also integrated with subsequent presentations, thereby producing a head start in the race process of choosing the more fluent word. If this prime-induced perceptual fluency favors the target, performance improves, but if it favors the foil, performance is harmed.

The effect of habituation through synaptic depression is twofold. First, habituation lessens the extent of residual activation. By itself, this only produces a reduction in the pattern of costs and benefits associated with priming. Critically, habituation also slows down the rate of processing in response to a repeated word, thereby causing a true perceptual deficit. This accounts for the full transition from integration to separation seen in the behavioral data as priming changes from positive to negative.

Although habituation exists at all levels within the model, it is habituation at the letter level that is most important for producing the transition from positive to negative priming. Thus, a "cognitive aftereffect" (Huber, 2008) for the letters of the prime makes it difficult to readily identify a repeat of those letters when they appear in a different screen location. The word level plays a role in producing repetition effects, but the word-level dynamics are too slow to fully explain the observed behavioral priming effects. Habituation at the lowest level cannot explain behavioral priming considering that this paradigm never repeats any words or letters in the same visual locations and so there is literally no difference between any of the priming conditions at the lowest level. However, as explained below, habituation of the low-level visual input is crucial for predicting electrophysiological effects.

A Priori Predictions for Electrophysiology Based on Behavioral Parameters

Considering that the target flash occurs 500 ms prior to initiation of the decision process (see the labels at the left-hand side of Figure 1 that indicate when event-related potential (ERP) and magnetoencephalography (MEG) responses are recorded compared to when behavioral responses are collected), early electrophysiological effects that modulate with prime duration would be suggestive of a perceptual account of priming rather than a decisional account. Because we examined brain responses within 200 ms after presentation of the flashed target and because there is no difference between the repeated and novel conditions until the target flash, explicit strategies cannot explain ERP/MEG repetition priming effects. Nevertheless, we do not deny the possibility that decisional effects can potentially explain response behavior even if there is also a perceptual effect. The question posed by this research is whether the observed perceptual effects are an adequate explanation of response behavior (i.e., whether decisional strategies are needed in addition, or whether perceptual effects are sufficient). To answer this question, we reverse engineer the problem, going from overt behavior back to perceptual responses. This was done by generating a priori predictions for perceptual ERPs from the previously developed neural habituation model with the previously published parameter (Huber & O'Reilly, 2003) that fit the data seen in Figure 2. Thus, behavioral results specified the time course of electrophysiological recordings and changes in these recording as a function of prime duration.

The parameters used to produce ERP predictions (see Figure 3B) for immediate word repetitions were found by fitting a particularly constraining data set that not only included the same conditions as Experiment 1 but additionally included other prime durations (as shown in Figure 2), other conditions with primes that were unrelated words (not shown in Figure 2), and "primes" that were simply pattern masks presented for equivalent durations (not shown in Figure 2). Huber and O'Reilly (2003) reported these

parameters and the associated simulation details and Huber (2008) reported the full experimental details and results. Therefore, these parameters were specified in advance based on behavioral data. Because nothing repeats at the lowest level of visual lines, the ERP predictions shown in Figure 3B are derived solely from the letter level (the yellow lines) and word level (the blue lines) as they sum to produce voltage potential differences between repeated and novel targets as recorded from posterior scalp locations (dashed black lines).

To fully understand these predictions, it is important to appreciate the difference between a field potential, which is the voltage potential associated with a large number of simultaneously active neurons, versus the response of a single cell. Electrophysiological recordings at the scalp are field potentials, and so they reflect the activation of entire cortical regions, such as the visual word form area (Dehaene, Le Clec, Poline, Le Bihan, & Cohen, 2002), rather than activations of single words. Therefore, predicted scalp ERPs are found by summing the membrane potential from all the simulated neurons within a level of the model (the yellow and blue lines of Figure 3B). Membrane potential is the most appropriate analogue of EEG for this model because scalp potentials are thought to arise from concurrent excitatory postsynaptic potentials at the apical dendrites of pyramidal cells (Kandel, Schwartz, & Jessell, 1991).

If the response to a word lingers for some period of time past its removal (i.e., integration), then the field potential to a subsequent word that is different than the prime (i.e., the novel condition) reflects not just the target word but also lingering activation from the prime word. Due to integration, presentation of a novel target word will produce a larger electrophysiological response compared to a repeated target word. Repeating a word may produce a larger response for the activation of the particular neurons associated with that target word, but the more important factor for field potentials is that presenting two different words involves a greater number of simultaneously active neurons and, thus, a greater summed response as measured at the scalp. Therefore, the general prediction is for 'repetition deficits' as revealed by smaller electrophysiological responses, although the term deficit is a misnomer because the underlying explanation is that two simultaneously active words add up to more than repeating a single word (i.e., a novelty boost is a more accurate description). These repetition deficits are apparent in Figure 3B at both the letter and word levels because the repeated conditions (dashed lines) lie below the novel conditions (solid lines), with the peak differences occurring in concert with peak activation.

Beyond this general prediction of repetition deficits, neural habituation predicts that these deficits will reduce in magnitude as a function of increasing prime duration. With a sufficiently long exposure to the prime, the perceptual input from the prime's location becomes habituated. Thus, activation of the prime word progressively diminishes at all levels with increasing prime duration. Even though the letter and word representations for the prime are likewise habituated, the dominant effect of habituation is a lessening of the perceptual input. When the same word reappears as the target in a different screen location, this provides a new burst of perceptual input (i.e., the prime and target are now separate events at the lowest level). In contrast, if the repetition occurs following a brief prime, the activation in response to the prime is still at peak values and the repeat produces little additional acti-

vation. In sum, the predicted interaction with increasing prime duration is that the repetition deficits (which are caused by two words summing to more than a repeat of one word) will diminish in magnitude. With separation due to sufficiently long prime exposure, any word, even a repeat, is viewed as a new perceptual event because the activation to the prime has sufficiently dropped off due to this low-level habituation. This is seen in Figure 3B by comparing the magnitude of the repetition deficit following a short, 150-ms prime (comparison of the solid and dashed lines in the top graphs) versus the magnitude of the repetition deficit following a long, 2,000-ms prime (comparison of the solid and dashed lines in the bottom graphs). This lessening of the repetition deficit with increasing prime duration is particularly notable for the letter level (yellow lines), whereas activation at the word level is slower to fade in response to low-level habituation and so there is less of a difference as a function of prime duration.

For ERP peaks, the P versus N letter designation refers to the polarity of the peak response (i.e., a positive versus a negative peak), and the number refers to the time in milliseconds at which that peak occurs in general. However, the underlying cortical sources of EEG are equivalent to electromagnetic dipoles (Kandel et al., 1991; Regan, 1989), and so for each P there should be a concurrent N on the opposite side of the head, although recording the opposite side of the dipole is often difficult or impossible if there are no electrodes at that location (e.g., in the neck). In the reported experiments, we examined the P100 and N170 peaks, which are measured with electrodes at posterior scalp locations in response to visually presented stimuli. In keeping with equivalent electromagnetic dipoles, our observations of P100s and N170s at posterior scalp locations were concurrent with N100s and P170s as measured with frontal electrodes, although we use the traditional terminology of P100s and N170s and we consider posterior electrodes as a scalp position grounding point.

Electrophysiological recordings not only reflect the summed response of entire brain regions (i.e., summation within levels of processing), but, in addition, these recordings reflect summation across different brain regions (i.e., summation across levels of processing). This form of summation across brain regions requires an anatomical model of where each region is located in relation to each electrode and the orientation of each brain region in its respective position. The neural habituation model specifies processing at each level at every millisecond, but it does not specify exactly where in the brain these regions are located or how they are oriented. Nevertheless, a rough approximation can be found by assuming that letter processing contributes to the P100 response and word processing contributes to the N170 response (in the General Discussion, we consider these assumptions in greater detail as they relate to the literature). Without any modification of the dynamics that best accounted for the behavioral data, the word level achieves its peak activation roughly 70 ms after the letter level. However, to accurately map these responses onto the P100 and N170, a temporal offset is needed between when a stimulus is first presented and when the line level of the model first receives input from that stimulus. If this offset is around 40 ms, which is roughly appropriate for the transduction time between the light signal and regions of visual cortex (Inui & Kakigi, 2006), then the times shown on the x-axis of Figure 3B correspond reasonably well to the P100 and N170.

With the simplifying assumptions that the source of the P100 is letter processing and that the source of the N170 is word processing and that the letter dipole is pointed towards the back of the head whereas the word dipole is pointed towards the front of the head, the dashed black lines at the bottom of the graphs in Figure 3B show the a priori predictions for ERP repetition effects. Under these assumptions, the predictions are found by taking the repeated minus novel responses at the letter level and subtracting this difference from the repeated minus novel responses at the word level. This second difference is the repetition effect solely based on the word level, and it is subtracted from the letter-level repetition effect because the word level is assumed to be oppositely oriented. The predictions for the time period corresponding to the P100 are relatively straightforward, with repetition deficits following both short- and long-duration primes, although the magnitude following a long-duration prime is predicted to be smaller due to habituation of the perceptual input. The predictions for the time period corresponding to the N170 are more complicated because there is a strong ongoing influence of the letter-level repetition effect during the N170 time period. In general, during the N170, the prediction is for greater responses with repetitions due to the opposite orientation of the word level (i.e., less of an N170 with repetitions corresponds to greater ERPs). However, rather than a lessening of the repetition benefit following a long-duration prime, the opposite is seen in Figure 3B, and the N170 achieves its greatest repetition benefit for the long prime duration condition. This is due to the overlapping nature of the two cortical sources—following a short-duration prime, there is still a large letter-level repetition deficit that partially counteracts the repetition benefit produced by the word level; but following a long-duration prime, this letter-level deficit is substantially reduced, which unmasks the underlying repetition benefit due to the word level.

One final aspect of these predictions concerns the difference between presenting just a single prime word (Experiment 1) versus presenting two different prime words (Experiment 2). The model includes inhibition within each level of processing, which serves to keep total activation from becoming excessive. Therefore, with two simultaneously active words, there is greater inhibition due to greater total activation, which tends to reduce the magnitude of the additional response when the target is presented. It is the letter level where this effect primarily occurs; and, as seen in Figure 3B, the differences between the repeated and novel conditions of the letter level are reduced for Experiment 2 compared to Experiment 1. As a result, the model predicts that P100 effects will be smaller in Experiment 2 and, paradoxically, that N170 effects will be stronger considering that the source of the N170 is not as strongly masked by the ongoing letter-level effect.

Outline of Experiments

Considering the complexities of reporting behavioral data, electrophysiological data, and a computational model that links the two types of data, we provide a brief preview of the reported studies. The overarching goal of this research was to test whether perception, rather than decision, might serve as a viable cause of the transition from positive to negative priming. This was tested by monitoring neural responses to the target word, which was presented prior to the choice words that initiate the decision process (see Figure 1). If the neural responses occur

sufficiently early as to indicate a perceptual effect and if the neural responses modulate as a function of prime duration, then these responses may serve as the basis of behavioral priming. Alternatively, if the neural sources underlying word perception provide a measurable signal, and yet there is a failure to observe repetition effects that modulate with prime duration despite sufficient power and repeated attempts (three experiments and two types of electrophysiological recording), this would cast considerable doubt on a perceptual account of priming.

The aim of these studies was to examine the nature of overlapping brain responses (i.e., integrated responses) when visual words are presented in rapid succession. Furthermore, by testing both short-duration primes and long-duration primes, we compared a situation that maximizes the entanglement of prime and target versus one were the prime and target were perceptually separated. We expected to see a complicated overlapping electrophysiological waveform following a short-duration prime, but it was literally this complicated overlap that we sought to study. We could have used any of a number of techniques that untangle overlapping waveforms through experimental manipulations, such as random variation in the timing of presentations. However, our goal was to directly study the entanglement and the manner in which the brain naturally untangles events as a function of processing time. We analyzed the results only in terms of repetition effects (i.e., differences between the repeated and novel conditions) so as to reduce the influence of other cortical regions that are not directly involved in the processing of visually presented words. Because the responses to the target were difficult to identify following a short-duration prime, we used the response to the prime in the long condition, which presented a single word in isolation, to determine appropriate time periods for the corresponding P100 and N170 responses to the briefly presented target word.

Two different ERP experiments were performed, with the first using a single prime word and the second two different prime words. Use of a single prime word produces maximal repetition effects (Experiment 1). Nevertheless, it could be that the difference between the short and long conditions is due to differential attention with the single prime paradigm. For instance, if attention is heightened 150 ms after presenting a word but then attention falls to a low level 2,000 ms after presenting a word, this difference in attention between short- and long-duration might explain the interaction between priming and prime duration. For this reason, the second experiment adopted a different procedure, one that eliminated this possibility by using visual presentations that were identical in all conditions prior to presentation of the target. In order to do this, it was required that two different prime words be presented on every trial such that the first prime served as the prime in the long condition and the second prime served as the prime in the short condition (see Figure 1). As explained above, this change in procedure was predicted to produce smaller P100 repetition effects. In keeping with these predictions, Experiment 1 produced both P100 and N170 repetition effects, whereas Experiments 2 and 3 only produced the later repetition effect (N170 for ERPs and M170 for MEG).

The a priori ERP repetition effect predictions provided an adequate qualitative explanation of the observed data. In order to assess the quantitative strength of these predictions and the extent to which the assumption of equal and oppositely oriented dipoles for the sources of the P100 and N170 was appropriate, we fit

anatomical dipoles for the letter and word levels to the data of each participant using the previously specified dynamics. Furthermore, we did this with the likelihood ratio chi-square goodness of fit statistic, which produces significance tests for the reliability of the fit. Unlike brain electrical source analysis (BESA; Scherg, 1990), which allows that dipoles take on different, independently determined magnitudes at each observed time, our equivalent dipole model is constrained to specific values across time. Because BESA is a popular method for extracting cortical source information (e.g., Bekker, Kenemans, & Verbaten, 2005; Qiu, Luo, Wang, Zhang, & Zhang, 2006; Reinvang, Magnussen, Greenlee, & Larsson, 1998; Thierry, Doyon, & Demonet, 1998), we directly compared our account to a temporally unconstrained dipole model. Therefore, we also fit a model that contained these same letter and word dipoles but allowed that each dipole take on any value at any moment in time, such as with a typical application of BESA. Thus, the neural habituation model was literally nested under this BESA model, and so the goodness of fit between each model was directly comparable in light of the differences in the number of free parameters through a chi-square difference test (e.g., Batchelder & Riefer, 1990; Collins, Fidler, Wugalter, & Long, 1993). This provides a comparison between these a priori dynamics of the habituation model versus the popular technique of using unconstrained dynamics. This allows us to statistically assess whether the specific dynamics of the habituation model are adequate compared to all possible other dynamics that the temporally unconstrained BESA model considers.

The results of dipole fitting revealed large differences across individuals for cortical location. Such differences might arise from anatomical differences or, instead, they might indicate the spatial limitations of scalp potentials. In order to discriminate between these possibilities and to provide replication of the electrophysiological repetition effects, Experiment 3 was identical to Experiment 2, but used MEG, which provides better spatial localization. The MEG results confirmed that the spatial layout of these early perceptual cortical responses varied substantially across individuals. In order to normalize against these individual differences, the response to the first prime word was used as a "standard" response to determine both the timing of the MEG responses, as well as the multivariate pattern of the response across all the sensors. The subsequent target responses were then compared to this standard response, producing a single number for the magnitude of the M100 (which is similar to the ERP P100) and the magnitude of the M170 (which is similar to the ERP N170). With these individual differences normalized both in spatial location and onset latency, the results were analyzed across individuals, replicating the findings of Experiment 2.

ERP Repetition Effects: Experiments 1 and 2

Using the experimental paradigm seen in Figure 1, participants performed forced-choice perceptual identification while voltage potentials were recorded from a 129-channel Geodesic Sensor Net (Electrical Geodesics Inc., Eugene, OR). The basic behavioral task was to identify the briefly flashed target word with a forced-choice decision. On half the trials, the prime repeated the correct target word (repeated condition) and on the other half of the trials, the prime repeated the incorrect foil word (novel condition). Furthermore, half the trials presented primes for 150 ms (short condition),

whereas the other half of the trials presented primes for 2,000 ms (long condition). Many aspects of the paradigm were designed to reduce reliance upon strategic decisional factors in relation to the primes. A prime duration of 150 ms was chosen for the short condition so that the primes in all conditions were above the threshold of awareness, thereby reducing concern for conscious control factors as the difference between short- and long-duration priming. Participants were fully informed of the design and told there could be no effective strategy in selecting for or against primed words, considering that the incorrect foil word repeated a prime just as often as it repeated the correct target word. To reinforce this assertion, trial-by-trial feedback was provided. ERPs were time locked to the onset of the target flash (i.e., before the choices were known), and the measure of interest was the difference between repeated and novel conditions.

Experiment 2 also examined responses to the target word following short- or long-duration primes, but this was done with two different primes on every trial (one long and one short), so that there was no difference between any of the conditions up until the time of the target presentation (the target might repeat the long-duration prime, the short-duration prime, or neither). This served two purposes. First, it controlled against the possibility that the prime duration effect of Experiment 1 was due to the presence of *any* briefly flashed words versus a prime word that remained onscreen for several seconds. A briefly flashed word may engage attentional and control factors that are different compared to a long-duration presentation. Therefore, Experiment 2 always included both a long-duration prime and a briefly flashed prime on every trial. Although Experiment 2 contained four conditions in terms of behavioral priming (as differentiated at the time of the choice words), there were just three conditions for the ERP results considering that all conditions were identical prior to the briefly flashed target. In other words, the briefly flashed target might be the same as the short prime, the same as the long prime, or unrelated to both prime words. Thus, this last condition served as the novel baseline condition for both the short- and long-duration repetition conditions. The other purpose of Experiment 2 was to test the prediction that presenting two different prime words would result in more activation, and therefore more inhibition, which would reduce P100 repetition effects (i.e., including extra prime words dilutes the letter-level effect of repeating a particular word).

Method

Participants

There were 33 participants in Experiment 1, and 34 in Experiment 2. They ranged in age from 18 to 21 years and all were right-handed native speakers of English with normal or corrected vision. Fifty of the participants were University of Colorado undergraduates receiving course credit for their participation, and the other 17 participants received monetary compensation. Two of the participants from Experiment 1 were not included in the ERP analyses due to insufficient numbers of valid EEG trials. Four of the participants in Experiment 2 were not included in any analyses. The reasons for these exclusions were a programming error resulting in a zero target flash duration for one participant, a participant who never understood the nature of the task, and two participants who forgot their glasses.

Materials

A pool of 1,000 five-letter words was used for all presentations in Experiment 1 and an expanded pool of 1,500 five-letter words was used in Experiment 2. These words had a minimum written language frequency of 4 per million for Experiment 1 and a minimum frequency of 2 for Experiment 2, as defined and measured by Kucera and Francis (1967). All words were displayed in upper case, Times Roman font, as gray lettering against a black background, and subtended approximately 1.6° of horizontal visual angle. Upper and lower primes (these were the same word presented twice in Experiment 1, but two different words in Experiment 2) were on two single-spaced lines, one above the other, and collectively subtended approximately 0.8° degrees of vertical visual angle. Stimulus materials were displayed on LCD monitors with presentation times synchronized to the vertical refresh. The refresh rate was 75Hz providing display increments of 13.3 ms.

Experimental Design

Experimental sessions lasted from 2 to 2.5 hours and included a practice session, a block of 64 trials during which appropriate target flash times were determined, and five blocks of 80 experimental trials. Between each block of trials, impedances were checked to makes sure all electrodes were below 50 kΩ. All variables were within subject. The basic design for both experiments used the following two variables: priming condition, with two levels (target primed or foil primed), and prime duration, with two levels (short or long). There were 100 repetitions of each of the four basic experimental conditions randomly intermixed across the five blocks of 80 trials. Word selection occurred randomly without replacement.

As seen in Figure 1, two prime words were presented on every trial, one immediately above and one immediately below the fixation point. In the case of Experiment 1, these two prime words were identical, both appearing for the same duration. In Experiment 2, there were two different prime words, with the top word appearing for 2,000 ms and the bottom word only appearing for the final 150 ms prior to the target flash presentation. During the forced-choice decision, the target and foil were presented to the left and right of fixation, with the left/right position of the target fully counterbalanced. On every trial, a sequence of events occurred as shown in Figure 1. Prior to the first display of the prime words, a blank screen appeared for 100 ms, followed by a row of five plus signs for 500 ms in order to fixate the participant's attention to the appropriate part of the screen. The forced-choice response was nonspeeded and error feedback ('correct' or 'incorrect') appeared for 500 ms prior to the start of the next trial.

Eight trials at a time were run during which participants were instructed to refrain from blinking or moving. At the end of each 8 trials, participants were allowed to relax until they felt ready to continue. There were 16 practice trials prior to the 64 target flash time determination trials. The target was presented for 187 ms during these practice trials. The 64 target flash duration trials began with a target flash time of 113 ms. At the end of every 16 trials, target flash times were adjusted in order to progressively reduce target flash duration until the perceptual threshold was obtained, with performance close to 75% correct. After these trials, target flash duration remained fixed for the remainder of the experiment.

EEG Methods

Scalp voltages were collected with a 129-channel Geodesic Sensor NetTM (Tucker, 1993) connected to an AC-coupled, 129-channel, high-input impedance amplifier (2,000 MOhm, Net AmpsTM). Amplified analog voltages (0.1–100 Hz bandpass, −3 dB) were digitized at 250 Hz. Individual sensors were adjusted until impedances were less than 50 kΩ. EEG was digitally low-pass filtered at 40 Hz. Trials were discarded from analyses if they contained eye movements (EOG over 70 μV), or more than 20% of channels were bad (average amplitude over 100 μV or transit amplitude over 50 μV). Trials were included regardless of accuracy. Individual bad channels were replaced on a trial-by-trial basis with a spherical spline algorithm (Srinivasan, Nunez, Tucker, Silberstein, & Cadusch, 1996). Consistently bad channels for a given participant were replaced throughout that participant's entire dataset. ERPs were created by averaging all acceptable trials for each participant in each condition, time locked to the target onset. The mean number of acceptable trials per subject and condition ranged from 94 to 96 for both experiments. EEG was measured with respect to a vertex reference (Cz), but an average-reference transformation was used to minimize the effects of reference-site activity and to accurately estimate the scalp topography of the measured electrical fields (Dien, 1998). Average-reference ERPs were computed for each channel as the voltage difference between that channel and the average of all channels. The average reference was calculated to correct for the polar average reference effect (Junghofer, Elbert, Tucker, & Braun, 1999). ERPs were baseline corrected with respect to a 100-ms pretarget recording interval.

Equivalent Dipole Methods

As explained in the model analysis section within the results, a dynamically varying strength of activation for each dipole was determined from the neural network behavior or through linear regression for the BESA model. Predicted scalp electrode voltage potentials were calculated from a linear summation of four equivalent dipoles, given their positions and dynamically weighted moments (Berg & Scherg, 1994). This was done assuming a spherical four-shell head model with the conductances reported by Stok (1986). Locations for the 129 electrodes were determined by projecting their standardized nonspherical locations onto the outer shell. Appropriate up-down and front-back offsets for this projection were calculated so as to minimize the square of the projection distances. The simulated scalp potentials to the novel condition were then subtracted from the simulated scalp potentials to the repeated condition and compared to the observed ERP differences (i.e., repetition effects).

Likelihood Ratio Chi-Square Goodness of Fit

There are three commonly used methods to assess goodness of fit. The first of these is a least squares measure, such as is used in regression analysis. However, least squares is purely descriptive and cannot be used for hypothesis testing beyond noting the percent variance accounted for. For hypothesis testing, there are

two different common methods for producing goodness of fit statistics that are distributed as a chi-square variable: Pearson chi-square and likelihood ratio chi-square (see Collins et al., 1993 for a discussion and comparison between these measures). There are reasons to prefer the likelihood ratio over Pearson chi-square (e.g., Gonzalez & Griffin, 2001; MacCallum, Browne, & Cai, 2006); but, beyond these reasons, real-valued ERP data are ill-suited to calculation of Pearson chi-square. Therefore, we derived the appropriate equation for calculating chi-square based on the likelihood ratio, which is traditionally termed G^2.

Derivation of G^2 requires calculation for the probability of an observed ERP value in reference to sampling noise (i.e., ERP standard deviation). In other words, even a "poor fit" in absolute magnitude is not necessarily a poor fit in terms of maximum likelihood if the observed data include a large degree of noise. Because the data of each participant were separately fit, we calculated ERP standard deviations separately for each individual, with resultant values ranging from 0.73 to 3.66 microvolts across individuals. In this manner, the quality of fit for each individual was calculated in light of the ERP standard deviation for that individual. This could have been done in terms of the standard deviation of each separate electrode for each individual, although doing so would introduce additional error because the calculation of ERP standard deviation would be based on a much smaller sample (i.e., the standard deviation could no longer be assumed to be a population value, which is a necessary assumption to calculate G^2). A measure of ERP standard deviation for each individual was found using the plus–minus technique (Handy, 2005), which reverses the mathematical sign for half of the trials within a condition before averaging. This eliminates the average signal and the resultant ERP waveform only reflects sampling noise. The standard deviation was then calculated across the 36 time slices of the P100 and N170 time periods and across the 129 electrodes, thus yielding a measure of variability based on 4,644 plus/minus ERP values per individual in each condition. Finally, the ERP standard deviations for each of the four conditions were averaged to yield one highly reliable measure of ERP standard deviation for each individual, which was assumed to be the population value for that individual. Under the assumption that ERP noise is normally distributed, these standard deviations were used to calculate the standard deviation of the ERP difference between the repeated and novel conditions (i.e., the variance of the difference is the sum of the variances).

The likelihood ratio chi-square goodness of fit statistic, G^2, states that -2 times the natural log of the likelihood ratio is distributed as a chi-square distribution with degrees of freedom equal to the difference in the number of free parameters between the model in the numerator and model in the denominator of the ratio (Hoel, 1971). In the current situation, the two models in the likelihood ratio are a constrained model with fewer parameters than observations versus an unconstrained model that perfectly fits the data (i.e., a model that is the data itself). In other words, if the ERPs are very noisy, data that perfectly match expectations may be a low probability event, and so a model that misfits the data is not viewed as excessively erroneous in light of the sampling noise. By assuming that ERP noise is normally distributed, the normal density function is used to calculate a likelihood ratio for each observation based on the difference between observed and predicted values for the constrained model versus no difference for the

unconstrained model (see Equation 1). Assuming that all observations are independent of each other, the likelihood ratio of the entire data set is the product of the likelihood ratios for each observation. This independence assumption follows from the assumption that the model is the "true model," thus implying that lack of fit only occurs due to sampling error. Finally, because G^2 is twice the log of the likelihood ratio, this measure of goodness of fit is easily calculated with the sum of the log likelihoods, which results in Equation 1, where N is the number of observations, σ_{ERP} is the ERP standard deviation (in the current situation, this is the standard deviation of the ERP difference between the repeated and novel conditions as explained above), and k is the number of free parameters in the model.

$$G^2 = -2\log\left[\frac{\prod_{i=1}^{N}\frac{1}{\sigma_{ERP}\sqrt{2\pi}}e^{-\left(\frac{(obs_i - pred_i)^2}{2\sigma_{ERP}^2}\right)}}{\prod_{i=1}^{N}\frac{1}{\sigma_{ERP}\sqrt{2\pi}}e^{-\left(\frac{0^2}{2\sigma_{ERP}^2}\right)}}\right]$$

$$= \frac{1}{\sigma_{ERP}^2}\sum_{i=1}^{N}(obs_i - pred_i)^2 \approx \chi^2_{df = N - k} \quad (1)$$

The numerator and denominator normal density values in Equation 1 are not technically probabilities. However, they could be converted to probabilities by assigning some small bin width to the normal distribution. Note that doing so would introduce the same bin width term to both numerator and denominator and these would cancel each other. So, in the end, calculation of a value that is distributed as a chi-square is quite simple if one assumes that ERP noise is normally distributed. The net result after algebraic simplification is just the sum of squared differences divided by ERP variance. The 15 free parameters were separately adjusted for each participant so as to minimize this chi-square for the 144 ms of the P100 and N170 time periods. With a sampling rate of 250 Hz, this corresponds to 36 separate time points across the 129 electrodes in the short and long conditions, and so 9,288 fitted data points per participant.

Finally, we note that the logic behind the likelihood ratio chi-square statistic is the same as the logic behind the general linear model that underlies the analysis of variance (ANOVA). As with a Model III general linear model for the ANOVA, in which the same residual error is applied for all conditions but different residual errors are sampled for different individuals, our use of G^2 assumes a constant ERP error at all time points but different ERP residual error for different individuals. This is not the same as assuming that different ERP components have the same variance; ERP component variance is typically calculated across several time points and across individuals who may differ in the magnitude or onset of the systematic aspect of the ERP waveform. The only difference between our application of G^2 and a model III ANOVA is that we have assumed population values for ERP residual error in light of the large sample sizes used to estimate the ERP residual error. This allowed us to use a chi-square test rather than an F test, which, in turn, allowed us to compare goodness of fit for the temporally constrained habituation model to a temporally unconstrained BESA model by means of the chi-squared

differences for nested models (see Batchelder & Riefer, 1990 for an example of nested model comparisons).

Results and Discussion

Below, we report three separate sets of results: (a) behavioral results, which replicate the behavior seen in Figure 2; (b) statistical analyses of the ERP data; and (c) dipole modeling results that compare performance of the neural habituation model to the best case scenario of a model that allows independent values at each time slice (i.e., a comparison to BESA with similarly constrained symmetric dipoles).

Behavioral Results

For Experiment 1, the average target flash duration was 50 ms, ranging from 27 ms to 80 ms across individuals. For Experiment 2, the average target flash duration was 97 ms, ranging from 23 ms to 190 ms across individuals. The procedure for setting target flash duration at an individual's threshold was not as effective in Experiment 2, resulting in high accuracy. Figure 4 shows probability correct for the four conditions of all three experiments. There was a main effect of priming condition in both experiments, $F(1, 32) = 15.5$, $p < .001$ and $F(1, 30) = 28.91$, $p < .001$. The main effect of prime duration was only significant in the first experiment, $F(1, 32) = 12.87$, $p < .01$, but not the second experiment, $F(1, 30) = 0.003$, $p = .96$. The interaction between priming condition and prime duration was of fundamental interest, and this was significant in both experiments, $F(1, 32) = 93.97$, $p < .001$ and $F(1, 30) = 82.48$, $p < .001$. For the short prime duration, performance in the target condition was better than performance in the foil condition in both experiments, $t(32) = 8.07$, $p < .001$ and $t(30) = 9.18$, $p < .001$. For the long prime duration, performance in the target condition was worse than performance in the foil condition

in the first experiment, $t(32) = 3.59$, $p < .001$, but not the second, $t(30) = 0.20$, $p = .84$.

Evoked Responses (ERPs)

We examined the earliest detectable ERP components in response to the briefly flashed target prior to the presentation of the choice alternatives. These components consisted of the first positive going wave, or P100 (as measured at posterior scalp locations), which peaked around 110 ms after presentation, and the first negative going wave, or N170, which peaked around 180 ms after presentation. As seen in Figures 5 and 6, these ERP components revealed significant repetition effects that were modulated by prime duration.

The ERP components in response to the brief target flash were small and overlapped greatly with the large ERP components elicited by the above-threshold prime presentations. Therefore, appropriate time periods for the target's P100 and N170 were determined from the prime's P100 and N170, which were more easily isolated. It is conceivable that the target's P100/N170 responses were offset in some systematic manner, considering that the target was only briefly flashed. However, the P100 to the briefly flashed target in Experiment 1 following a long-duration prime was easily isolated, as seen in Figure 5, and confirmed the offset technique; this P100 occurred at the expected time based upon the P100 to the long-duration prime, thus validating use of the long-duration prime to determine appropriate time windows. The prime's P100 and N170 were separated by approximately 70 ms in both experiments, although these responses were advanced by 20 ms in Experiment 2.

In analyzing these data, the effect of prime duration on repetition priming was of critical importance. The neural habituation model predicted that the magnitude of the repetition effect should dimin-

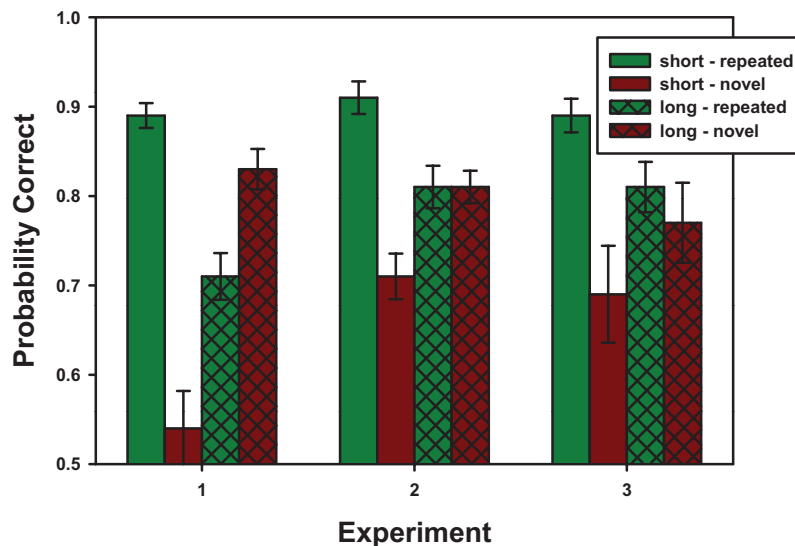*Figure 4.* Behavioral data (probability correct with forced-choice testing) for the three reported experiments. Error bars are plus and minus one standard error of the mean. The legend refers to the prime conditions. Experiment 1 used a single prime word and produced both larger positive priming effects following short-duration primes as well as larger negative priming effects following long-duration primes.
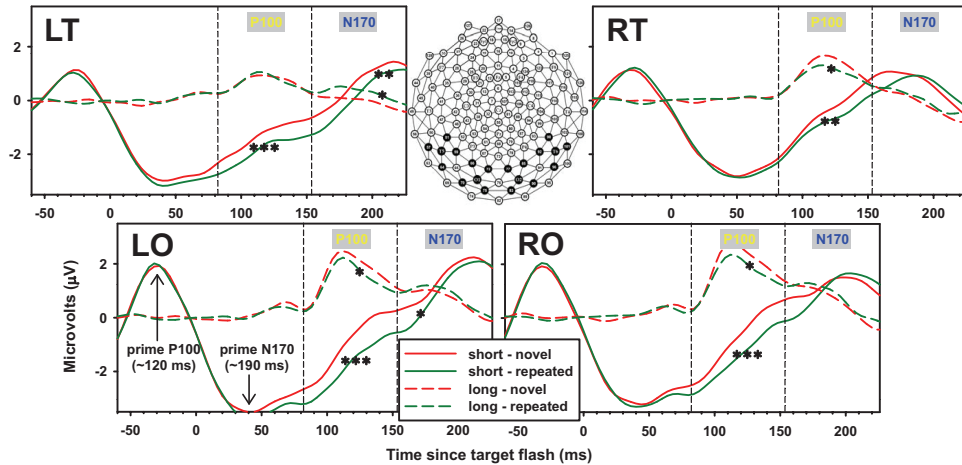
Figure 5. Average event-related potentials (ERPs) for Experiment 1, which used a single prime word presented for 150 ms (short) or 2,000 ms (long). The graphs show results for the four groups of electrodes shown in black on the electrode map. There were five electrodes per group, and groups were labeled left temporal (LT), left occipital (LO), right occipital (RO), and right temporal (RT). The legend refers to the prime conditions. The stars indicate different significance levels for the ERP repetition effects. $^* p < .05.$ $^{**} p < .01.$ $^{***} p < .001.$

ish with increased prime duration. Furthermore, assuming that the letter level is the source of the P100 and the word level is the source of the N170, with these levels equal in magnitude, the model predicted P100 repetition deficits, with these deficits decreasing in magnitude as a function of prime duration, as well as N170 benefits, with these benefits growing in magnitude as a function of prime duration. Finally, the model predicted that the P100 effects in Experiment 2 would be smaller due to use of two prime words.

To assess the reliability of the observed perceptual ERPs, we statistically analyzed 20 posterior electrodes. Recent analyses of EEG data with independent components analysis of the visual P100 and N170 suggests that limiting consideration to posterior electrodes minimizes the effect of more anterior sources (Makeig et al., 2002). To assess topographic differences, these electrodes were further divided into four groups with five electrodes per group (see Figures 5 and 6), arranged over both hemispheres, in either occipital or temporal scalp regions. We term these electrode
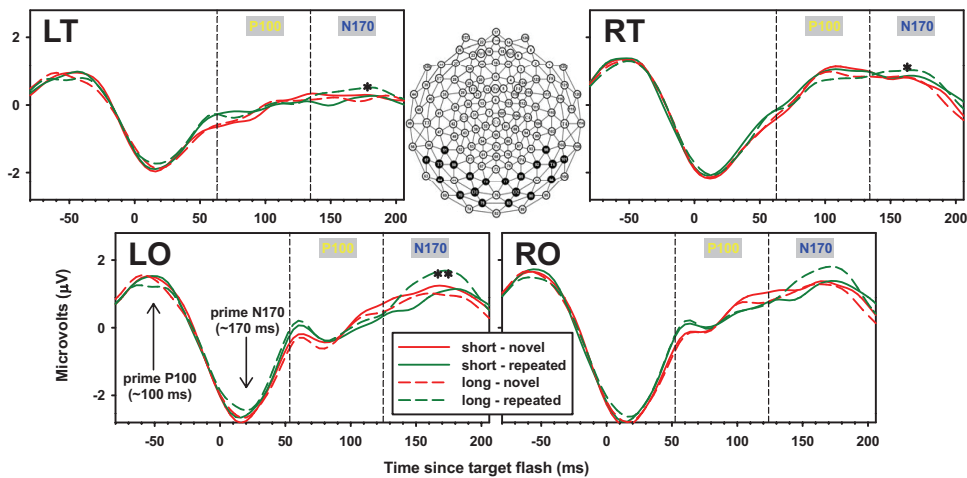
Figure 6. Average event-related potentials (ERPs) for Experiment 2, which used two different prime words with the top prime presented for 2,000 ms and the bottom for the final 150 ms such that all conditions were identical up until the time of the target flash. The legend refers to the prime conditions. See the Figure 5 caption for other descriptions. In terms of these ERP responses, which occur prior to the choice words, the short-novel and long-novel conditions are in fact the same condition at this point in the presentation sequence for the design of Experiment 2, although the graphs show them separately, providing an indication of sampling error. $^* p < .05.$ $^{**} p < .01.$

groups left occipital (LO), right occipital (RO), left temporal (LT) and right temporal (RT). Repeated measures ANOVAs were run on the mean amplitude of the repetition priming difference scores (i.e., ERP repetition effects), which were averaged across the eighteen 4-ms time slices comprising the separate time periods where the P100 and N170 responses to the target flash were expected to be maximal. For Experiment 1, the P100 time period was the average response from 84 ms to 152 ms after the target flash, and the N170 time period was the average response from 156 ms to 224 ms (see Figure 5). For Experiment 2, the P100 time period was the average response from 64 ms to 132 ms after the target flash, and the N170 time period was the average response from 136 ms to 204 ms (see Figure 6). The included variables were hemisphere (left, right), region (occipital, temporal), and prime duration (short, long). In addition, Figures 5 and 6 label all significant ERP repetition effects for each electrode group.

First, we consider the N170 time period, which was more similar across the experiments. There was a main effect of prime duration on the N170 repetition effect for both experiments, $F(1, 30) = 6.00$, $p < .025$, and $F(1, 29) = 4.31$, $p < .05$. For Experiment 1, there was a two-way interaction between duration and hemisphere, $F(1, 30) = 5.15$, $p < .05$, as well as a three-way interaction between duration, hemisphere, and region, $F(1, 30) = 4.93$, $p < .05$. Specific contrasts identified these effects as centered in the LT group, which changed from a negative ERP repetition effect following a short-duration prime, $t(30) = 3.05$, $p < .01$, to a positive ERP repetition effect following a long-duration prime, $t(30) = 2.40$, $p < .05$, whereas the LO group changed from a negative ERP repetition effect, $t(30) = 2.31$, $p < .05$, to an absent effect, $t(30) = 0.77$, $p = .45$. There were no significant repetition effects for the RT and RO groups. For Experiment 2, there were no significant two-way or three-way interactions. The basic result for Experiment 2 was a positive repetition effect following a long-duration prime, $t(29) = 2.84$, $p < .01$, but no effect following a short-duration prime, $t(29) = 0.31$, $p = .76$.

Summarizing across the relatively minor regional differences, the N170 to a repeated word (compared to a novel word) was larger following a long-duration prime (both experiments), but was either smaller (Experiment 1) or equivalent (Experiment 2) following a short-duration prime.

Next, we consider the P100 time period, which differed more greatly between the experiments, as expected, due to the difference between using one versus two prime words. For Experiment 1 there was an effect of duration, $F(1, 30) = 10.65$, $p < .01$, and region, $F(1, 30) = 12.90$, $p < .01$ for the ERP repetition effects, but no interaction. The t tests against a test value of zero revealed a negative ERP repetition effect both following short, $t(30) = 5.66$, $p < .001$, and long prime durations, $t(30) = 1.83$, $p < .05$, although this effect was larger for short prime durations $t(30) = 3.26$, $p < .01$. Furthermore, the ERP repetition effect was larger for the bilateral occipital electrode groups than the bilateral temporal groups, $t(30) = 3.59$, $p < .01$. In contrast, for Experiment 2 there were no significant effects of duration, hemisphere, or region, but there was a significant interaction between duration and region, $F(1, 29) = 5.69$, $p < .05$. Although there were no significant ERP repetition effects for any of the electrode groups, there was a tendency for the ERP repetition effect to decrease for the bilateral occipital groups as a function of prime duration, whereas there was a tendency for the ERP repetition effect to increase for

the bilateral temporal groups as a function of prime duration. The significant difference between these two tendencies, $t(29) = 2.39$, $p < .05$, resulted in the observed interaction between prime duration and electrode region.

In summary, there were P100 deficits to repeated words, and these deficits diminished with prime duration and occurred more at occipital scalp locations. In the case of Experiment 1, there was a significant negative ERP repetition effect at both prime durations, whereas in Experiment 2 there were no significant ERP repetition effects, although there was a tendency for the slightly negative ERP repetition effect to diminish with prime duration (which contributed to the significant interaction).

In summary, both P100s and N170s revealed perceptual repetition effects in response to the target word, and these effects varied as a function of prime duration. These findings are problematic for the claim that discounting occurs in response to the choice words rather than the briefly flashed target word. Furthermore the qualitative pattern of these perceptual repetition effects was as predicted a priori by the neural habituation model: (a) diminishing P100 repetition deficits with increasing prime duration; (b) increasing N170 repetition benefits with increasing prime duration; and (c) smaller P100 effects with use of two primes rather than one prime. The only aspect that did not conform to prediction was the appearance of N170 repetition deficits in the short condition of Experiment 1. However, it is important to keep in mind that these predictions were based on exactly equal and opposite contributions for the letter and word levels. If it is assumed that the letter level produced an overall slightly larger response than the word level for the participants of Experiment 1, then the model can quantitatively capture this crossover interaction as seen next. Essentially, with a larger letter-level response, there was sufficient masking from the ongoing P100 deficit during the N170 time period following a short-duration prime as to produce the appearance of N170 repetition deficits. Only following a long-duration prime was the underlying N170 repetition benefit revealed.

Modeling Analyses

Because the neural network only includes the minimal representation needed for word identification, irrespective of task demands, attention, and other cognitive functions, it would be inappropriate to apply it to the ERP waveforms of a single condition, particularly in the case of the short conditions in which later-acting, higher order responses to the primes (e.g., P300 and N400) overlapped with the target ERP response. However, by limiting the analysis to the difference between conditions (i.e., repetition effects), the model provides a remarkably accurate account of the ERP repetition effects, both as a function of time and scalp location.

The fundamental prediction of the model is for less of a response for a repeated item due to integration. However, this effect should be diminished following a long-duration prime (see Figure 3B and the associated discussion). In mapping this prediction onto scalp potentials, it is necessary to specify the position and orientation of the underlying cortical sources (see Figure 3C for the best fit average dipole positions and orientations). This was achieved by quantitatively fitting ERP repetition effects at every electrode and every 4-ms time slice during the P100 and N170 time windows using a unique equivalent dipole methodology in which dipoles

provide a static position and a time-invariant magnitude that is multiplied by the dynamically varying neural network activations to produce dipole activation values. The reported dipole position fits were done separately for the data of each participant. Unlike BESA dipole modeling (Scherg, 1990), which allows that dipoles can be weighted by a different, independently determined value at every moment in time, this solution is highly constrained, employing fixed dynamics as specified by the previously reported fit of the neural network to the behavioral data (i.e., the a priori dynamics seen in Figure 3B are not allowed to vary).

Parameter estimation and nested modeling. In producing an account of the ERP data, it was assumed that the summed membrane potential of all units in a given level of the neural network produce a regional potential that can be simulated with an equivalent electromagnetic dipole. The spatial distribution of the scalp potentials depends upon the position and orientation of these equivalent dipoles, with position and orientation dictated by the geometry of the cortical folds (Kandel et al., 1991) and layer of cortical enervation (Regan, 1989). The only parameters allowed to vary were dipole positions (x, y, and z), dipole orientation (vertical and horizontal angle in spherical coordinates), static dipole magnitude, and the time offset between presentation of a word and the first effect of that word at the lowest level of the model. Dipoles were constrained to be symmetric across the hemispheres (i.e., no additional position or orientation parameters were needed for the opposite hemisphere dipole), although a different static magnitude was used on each side to capture hemispheric differences (i.e., each level of the word network was allowed to exist in each hemisphere to some extent). The first level of the neural model was simulated but did not specify a dipole because it is equivalent in its response in all conditions (i.e., as an ERP source, the first level contributes nothing to the difference between conditions). Therefore, four dipoles were used in fitting the data (one for letters and one for words in each hemisphere), although position and orientation parameters were needed for only two of the four dipoles considering the constraint of hemispheric mirror symmetry. This method of fitting the a priori network dynamics to the scalp potentials was run separately for the data of each individual for the 129 ERP data points collected in the short and long conditions over the 36 time slices, yielding 9,288 fitted values for 15 free parameters. The model results were then averaged across participants to produce the displayed ERP repetition effects. Parameters were optimized using chi-square error (Equation 1).

Chi-square values for the goodness of fit for the neural habituation model provide a probability measure of the fit. However, as pointed out by Roberts and Pashler (2000), there is more to evaluating a model than simply demonstrating that a fit is quantitatively good; it is important to engage in model comparisons and to consider model flexibility by examining what types of data a model cannot address. Therefore, we performed a second fit of the data using traditional equivalent dipole methodology, similar to the BESA method (Scherg, 1990), which is widely used to provide cortical position as well as response magnitude for assumed underlying neural sources (e.g., Bekker et al., 2005; Qiu et al., 2006; Reinvang et al., 1998; Thierry et al., 1998). This can be viewed as the best case scenario against which the fit with a priori neural habituation dynamics is compared. BESA methodology does not temporally constrain dipole responses, allowing uncorrelated responses from moment to moment and even reversals in the direction of response (Berg & Scherg, 1994). Producing dipole fits with these relaxed assumptions allowed comparison of the a priori dynamics against freely varying dynamics that are optimized to capture the data.

Unlike a typical implementation of BESA, we constrained the dipoles of each hemisphere to be mirror symmetric and of equivalent response at each moment, but with a free parameter for the weighting of one side versus the other. With these constraints on the BESA model, the neural habituation model is literally a subset of the possible time-independent equivalent dipole models and is thus nested under the BESA fit (i.e., the BESA fit can fit anything that the neural habituation model can fit, plus other results that the neural habituation model cannot fit due to its assumption of independence over time). This allows nested model comparisons in terms of the relative decrease in chi-square goodness of fit in light of the number of extra parameters contained in the more flexible parent model (Batchelder & Riefer, 1990; Huber, 2006). A nested model comparison is particularly desirable because it allows for statistical comparison without need of more complex techniques that attempt to calculate the separate flexibility assigned to each parameter when comparing nonnested models (Pitt, Kim, & Myung, 2003; Pitt & Myung, 2002; Pitt, Myung, & Zhang, 2002). Because equivalent dipoles sum in a linear fashion to produce scalp potentials, the separate values for each dipole at each moment are easily calculated using multiple linear regression. For the BESA fit, a static magnitude for each hemisphere was no longer needed, although a single weighting value was still needed to capture the relative contribution of the hemispheres. Therefore, the 14 static free parameters were reduced to 12 free parameters. Instead of a static magnitude multiplied by a priori dynamics, a different weighting value was needed for each of these two dipoles at each of the 36 time slices in both the long and short conditions, and so 156 free parameters were needed to fit the 9,288 data points collected for each individual; even though these parameters are found through multiple linear regression, they are properly viewed as free parameters and increase the flexibility of the model. Thus, there were more than 10 times as many free parameters for the BESA fit, with these extra parameters providing extra freedom over time to the same number of symmetric dipoles.

Temporal results: model comparison. Next we describe the results of model fitting both for the habituation model, which uses a fixed dynamic time course, as well as for the temporally unconstrained BESA model, which considers all possible time courses but uses the same set of mirror symmetric dipoles as the habituation model. A statistical comparison between these two models allows us to ask whether the prespecified dynamics of the habituation model are adequate in capturing these data.

In addition to the results at the level of individual subject fits, the model behavior for each individual was averaged and compared to the observed data collapsed across subjects. For this reconstituted subject-averaged data, the results across the 36 time slices for the a priori neural habituation dynamics are shown in Figure 7 both as a scatter plot over the 9,288 values (upper graphs) and at each moment in time for each experiment and condition averaged across the 129 electrodes (lower graphs). For the lower graphs, a straight average of the 129 electrode ERP repetition effect difference scores at a particular moment in time would produce a zero value, considering that frontal scalp electrodes show opposite responses to posterior scalp electrodes (and because
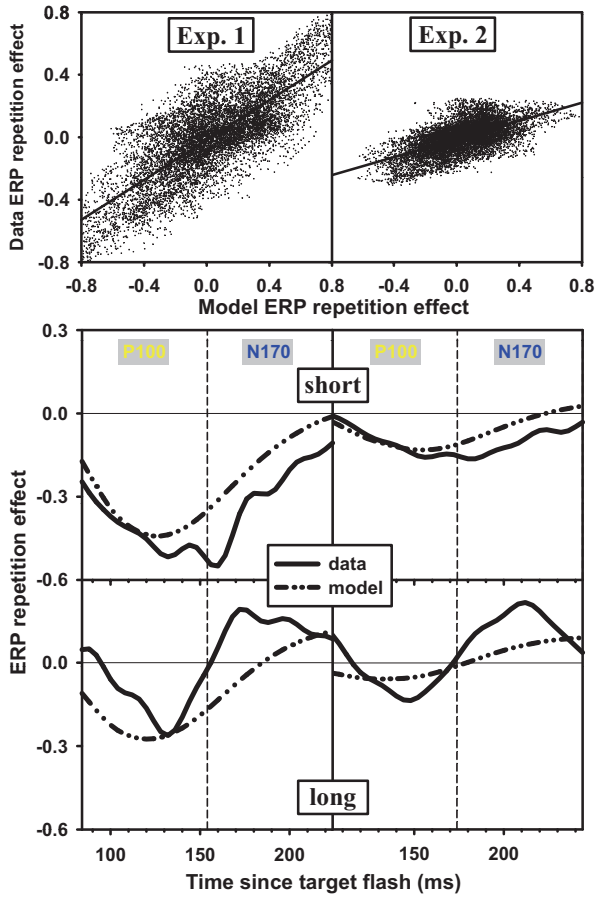
Figure 7. Temporal results from the habituation model for the subject-averaged data for Experiment 1 (Exp. 1) and Experiment 2 (Exp. 2). The upper scatter plots show the 9,288 event-related potential (ERP) repetition differences for the thirty-six 4-ms time slices of the P100 and N170 in the short and long conditions for all 129 electrodes. Goodness of fit regression lines are shown. The lower graphs show the average values at each time slice (see text for details of averaging technique). These graphs are directly comparable to the a priori predictions shown in Figure 3B, except that, in this case, dipole positions and orientations were fitted to the data (see Figure 3C) rather than assumed to be exactly equal and opposite in value.

an average reference was used). Therefore, the values were adjusted before averaging in producing Figure 7 such that electrodes producing positive ERP repetition effects during the P100 time window were flipped in numerical sign to convert them to negative ERP repetition effects. In other words, both the observed and model values for each electrode were multiplied by $+1$ or -1 depending on the average observed ERP repetition effect during the 18 time slices of the P100 time window. Essentially, frontal electrodes were converted to posterior electrodes before averaging across electrodes.

As seen in the lower graphs of Figure 7, the fixed neural habituation dynamics provide a reasonable account of the ERP repetition effects at each moment in time across the P100 and N170 time periods (more so for Experiment 1 than for Experiment 2). The dynamics for the a priori predictions seen in Figure 3B are the same, and so Figure 7 is necessarily very similar to 3B except that the values are shifted downwards slightly because dipole fitting placed larger static weights on the letter dipoles (the source of the P100) compared to the word dipoles (the source of the N170). Figure 7 does not show the best case scenario results of the time independent traditional dipole model; but, if it did, the graphs would reveal nearly perfect fits, with BESA capturing every seemingly random bump and wiggle at each moment in time. Thus, it is nearly certain that the assumption of independence over time results in fitting noise, which is why a quantitative comparison based on chi-square is needed to appropriately factor in the role of model flexibility.

First, we consider chi-squared error, which can be converted into a probability for the goodness of fit in light of the number of fitted values, the number of free parameters, and sampling noise. Chi-squared error was calculated according to Equation 1 (see associated discussion for calculation of sampling noise) for the data of each individual, with average results under both models shown on the left side of Table 1 in the last row. To calculate chi-squared error for the reconstituted group data (shown on the right side of Table 1), the ERP standard deviations found through the plus/minus technique for each participant were averaged and then divided by the square root of the number of subjects to provide a measure that appropriately reduced the expected sampling variance in light of the larger number of data points included in the group data. Table 1 does not report the associated probabilities for goodness of fit because all models fit excessively well, producing *p* values of exactly 1.0 (i.e., the probability that the

Table 1
Goodness of Fit Statistics for ERP Repetition Effects

Model	Individual data (9,288 values)				Group data based on fits to individual data (9,288 average values)			
	Experiment 1		Experiment 2		Experiment 1		Experiment 2	
	Habituation	BESA	Habituation	BESA	Habituation	BESA	Habituation	BESA
Number of parameters	15	156	15	156	465	4,836	450	4,680
Squared error (microvolts)	0.81	0.50	0.74	0.51	0.05	0.02	0.03	0.02
Variance accounted for	29%	54%	26%	48%	65%	86%	31%	58%
Correct valence	71%	77%	69%	75%	76%	86%	73%	79%
Chi-squared error	2,884	1,771	2,403	1,645	4,342	1,796	2,133	1,308

ERP = event-related potential; BESA = brain electrical source analysis.

model is rejected is 0.0 in all cases). Thus, the neural habituation model, with fixed dynamics and only 15 parameters successfully captured the 9,288 values at the level of individual data and at the level of the reconstituted subject-averaged data.

Providing other measures of model performance, Table 1 also reports average squared microvolt difference per electrode, percent variance accounted for, and a qualitative test for the proportion of data points that were of the correct numerical sign (i.e., correct valence—both model and data revealing positive or negative repetition effects). Even though the variance accounted for measure suggests a mediocre fit, it is important to keep in mind that (a) the variability in the ERP data is quite high and the model is doing remarkably well in light of the number of free parameters compared to the number of observed values and (b) the BESA fit, with more than 10 times as many parameters, is not doing much better in terms of these measures.

Neither the habituation dynamics nor the time impendent BESA model can be rejected based on goodness of fit. However, because this is a nested model comparison, the difference in chi-squared error between these models can be tested with degree of freedom equal to the difference in the number of free parameters. This asks the question of whether the extra parameters of the time independent BESA model produce a significantly better fit.[1] The answer to this question is quite interesting and depends on whether one examines the individual data or the group-averaged data. In terms of fits to individual data, the extra flexibility associated with the time-independent BESA model is warranted for most participants (although in a couple of cases, model comparison preferred the more constrained habituation dynamics). However, for the chi-square calculated for the average data based on the individual fits, the habituation dynamics are the statistically preferred model, and the extra parameters of the BESA model do not produce a significantly better fit, Experiment 1: $\chi^2(4,371, N = 9,288) = 2,546$, $p = 1.0$, Experiment 2: $\chi^2(4,230, N = 9,288) = 825$, $p = 1.0$. This is a sensible result considering that these particular a priori habituation dynamics were determined based on fits to group-averaged behavioral data (Huber & O'Reilly, 2003). In other words, the habituation dynamics of the *average participant* are a more parsimonious account of the ERP repetition effect of the average participant. Perhaps habituation dynamics fit to the behavioral data of each individual would fare better in terms of the ERP dynamics of each individual. In keeping with this conclusion, we recently performed a series of priming experiments with face perception using the same design as Experiment 1 in which we found reliable individual differences in the rate of prime induced habituation (Rieth & Huber, 2008).

Spatial results: Dipole positions and orientations. The previous section and results in Figure 7 emphasized the results in terms of each 4-ms time slice, whereas the current section examines the results in terms of topographic patterns across the 129 scalp electrodes for the P100 and N170. Figure 8 shows topographic maps for two time windows 72 ms apart, chosen to highlight the P100 and N170 spatial patterns (these are chosen sometime past the peak response so as to minimize the overlap of the P100 pattern during the N170). In this figure, the first and third rows show the observed ERP repetition effects, with dots indicating electrode position on this flattened scalp representation. The second and fourth rows show the corresponding spatial patterns produced by the a priori habituation dynamics after fitting the hemisphere
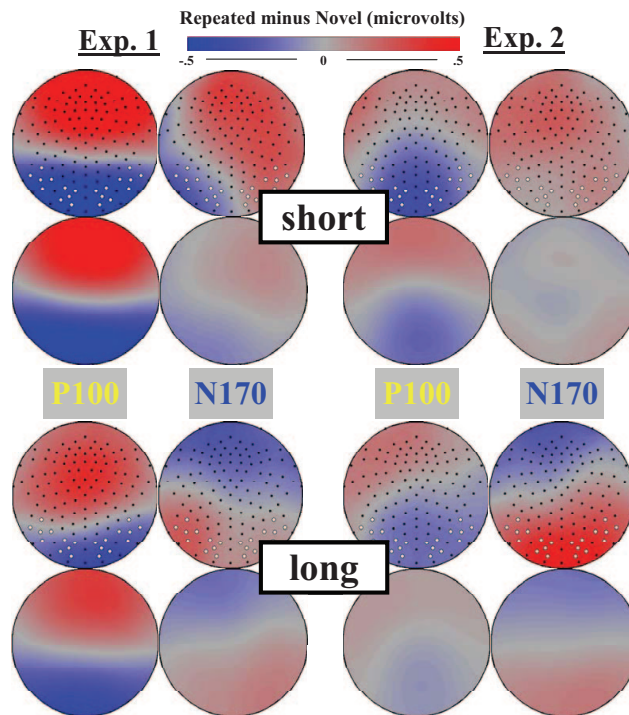

Figure 8. Spatial results from the habituation model for the subject-averaged data. Observed data (first and third rows with electrode positions indicated by dots) and neural habituation dipole fitted data (second and fourth rows) are broken down by prime duration (short, long). Two time slices 72 ms apart are shown to highlight the P100 and N170 spatial patterns. These occurred at 132 ms and 204 ms for Experiment 1 (Exp. 1) and 152 ms and 224 ms for Experiment 2 (Exp. 2). White dots indicate the electrodes used for statistical analyses.

symmetric dipole positions and orientations. Qualitatively, the spatial patterns appear close to the observed results. A more accurate quantitative fit could be found through the inclusion of more dipoles or by relaxing the hemispheric symmetry constraint, although doing so necessarily increases model flexibility and the number of free parameters.

Because dipoles were determined separately for the data of each participant, standard deviations can be calculated for the resultant parameters. By fitting the data of different individuals, which are subsets of the grand average data, this is similar to a nonparametric

---

[1] Fits with 4-ms time slices may appear biased in favor of the habituation model, considering that the number of free parameters in the time-independent BESA model depends on the chosen temporal resolution. However, this is not necessarily the case. Although a less fine-grained temporal resolution would entail a smaller difference in degrees of freedom due to fewer parameters in the BESA model, it would also entail fewer observations that need to be fit and would involve fitting less noisy data due to time averaging. The assumption of time independence in the BESA model allows it to capture truly random fluctuations from moment to moment, and so, to the extent that the fitted data are noisy and include such fluctuations (e.g., such as with small time slices), BESA will perform relatively better on average for each observation. The particular choice of 4-ms time slices is in keeping with the literature and typical use of BESA and is appropriate to the time scale of change in the habituation model.

bootstrap analysis for determining parameter confidence limits of the group data (e.g., Efron & Tibshirani, 1993), the only difference being that an analysis of individual data results in larger parameter confidence limits because it appropriately includes both random error as well as systematic subject variance. The resultant means and standard deviations across subjects of the model parameters are found in Table 2. As seen in this table, the standard deviations were large and similar in magnitude to the mean parameter values (this was equally true for the time independent BESA fit of the data). With such large deviations for the parameters, the resultant average dipole positions are not particularly reliable, and it would be premature to speculate on the cortical locations of these dipoles.

Relationship to other ERP analysis techniques. Similar to the adjacent response (ADJAR) technique for separating the components of overlapping waveforms through mathematical deconvolution based on uncontaminated waveforms (Woldorff, 1993), our analyses used the uncontaminated prime word ERP to determine appropriate times for the P100s and N170s to the target word combined with the neural habituation model to specify the nature of the separate overlapping components. However, our theory supposes that uncontaminated waveforms are already composites of multiple sources (e.g., the N170 partially reflects the ongoing response of the cortical area primarily responsible for the P100). Therefore, according to our theory, ADJAR deconvolution based on waveforms still fails to fully deconvolve the waveform into separate cortical responses, which might separately vary with the conditions of interest. This is the basic difficulty of the so-called "inverse problem" of electrophysiology (Mosher, Baillet, & Leahy, 1999), which arises because there is no unique solution for the mapping between an array of passively monitored three-dimensional sensors and the underlying cortical responses. To address this problem, additional assumptions are needed; and, in the current situation, the assumption we make is based on neural habituation as specified by the behavioral results.

There are other promising techniques that can similarly demix scalp responses, separating out concurrently active cortical sources by means of statistical covariance (Baillet & Garnero, 1997; Makeig, Jung, Bell, Ghahremani, & Sejnowski, 1997). A serious limitation of these techniques is the assumption of independence between the underlying sources (i.e., if one source is strongly active on a given trial, this does not imply anything regarding the activity of other sources). For removing eye blink artifacts or heartbeats, the independence assumption appears valid (Jung et al., 2000), but for differentiating between levels of perceptual processing, the independence assumption is almost certainly false (Halgren et al., 1995).

MEG Repetition Effects: Experiment 3

Experiments 1 and 2 obtained neural repetition effects for immediately repeated words, with these effects occurring in early perceptual responses. Furthermore, the magnitude of these responses varied with prime duration similar to behavioral priming. Because these were in response to the briefly flashed target word, prior to initiation of the decision process with the appearance of the choice words, this suggests a perceptual basis for the changes in priming as a function of prime duration. The combination of equivalent dipoles and the behaviorally specified neural network dynamics provided a reasonable and statistically reliable account of the electrode repetition effects seen in Figures 7 and 8, and this provides converging evidence to the claim that perception rather than decision underlies the effect of prime duration. In doing so, the complex data pattern was explained in terms of temporal and spatial overlap between the cortical source of the P100 and the cortical source of the N170. Furthermore, the difference between the two experiments was predicted due to the lessening of the repetition effect for the cortical source of the P100 when using two prime words rather than a single prime word. However, one troubling aspect of the modeling was that it failed to produce consistent cortical locations for the dipoles across the separate fits to the data of different individuals (see Table 2 and the associated discussion). This variability across individuals might indicate that the ERP data are spatially unreliable, perhaps because average electrode positions were used rather than determining exact electrode position for each participant or, perhaps, due to sampling noise. Alternatively, if these differences are real, this variability may indicate that different cortical areas are involved for different people or that different distributions of cortical areas are involved,

Table 2
Average Parameters and Standard Deviations

| | Experiment 1 | | | | Experiment 2 | | | |
| | Letters | | Words | | Letters | | Words | |
Statistic	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Left–right position	−0.14	0.14	−0.12	0.14	−0.14	0.16	−0.15	0.13
Front–back position	−0.02	0.32	−0.04	0.35	0.04	0.38	0.06	0.34
Up–down position	−0.17	0.22	−0.05	0.26	−0.11	0.24	−0.15	0.23
Horizontal angle	−0.84	1.59	0.60	1.62	−0.41	1.81	0.38	1.57
Vertical angle	−0.11	0.44	0.06	0.34	0.02	0.56	0.00	0.63
Response magnitude[a]	58.61	73.89	47.84	63.19	43.71	65.86	56.44	135.94
Right/left balance[b]	0.51	0.25	0.54	0.28	0.51	0.37	0.49	0.37
Time offset[a]	83.20	41.12			123.93	51.00		

[a] For the BESA fit of the data, these parameters were not needed, and, instead, a multiple linear regression across the two combined left/right dipoles was performed at each of the thirty-six 4-ms time slices for the short and long conditions, which added 144 more parameters instead of these three parameters.
[b] Parameters are for the left hemisphere dipoles, with the right hemisphere dipoles mirror symmetric and scaled by right/left balance.

with the dipole position indicating the central tendency of the distribution (i.e., each dipole was capturing a mixture of multiple cortical areas).

Experiment 3 replicates Experiment 2 using magnetoencephalography (MEG), serving two goals. First, replication helps establish that the effects seen Experiment 2 are reliable and generalize to other measures of neural activity. Second, MEG is more spatially localized in its response to cortical activity (Hamalainen, Hari, Ilmoniemi, Knuutila, & Lounasmaa, 1993), allowing a more reliable measurement of dipole position. Even though the use of two primes revealed weaker effects, the design of Experiment 2 is chosen for this replication because it controls against attentional differences between short- and long-duration primes by eliminating any differences between conditions prior to the target presentation. By convention, evoked MEG components are given the label M, indicating magnetic, followed by the number in milliseconds at which these responses typically occur poststimulus. The magnetic field responses are signed in terms of sinks (into the head) versus sources (out of the head), but the sign is not included in the labeling, presumably because MEG typically reveals both sinks and sources in close proximity due to its increased spatial localization. Therefore, the equivalent of the P100 ERP response is the M100 MEG component and the equivalent of the N170 ERP response is the M170 MEG component (Hopf, Vogel, Woodman, Heinze, & Luck, 2002).

The excitatory responses on the apical dendrites of pyramidal cells are the underlying source of both EEG and MEG (Regan, 1989), producing voltage potential changes as measured with EEG as well as magnetic field fluctuations as measured with MEG. Therefore, both techniques have the same high degree of temporal resolution. Nonetheless, there are some differences between the techniques. According to the right-hand rule of electromagnetism, a voltage potential and the associated current in a particular direction produces a magnetic field that wraps around the current dipole. Apical dendrites are aligned orthogonal to the cortical surface and so EEG reveals the largest response on the scalp directly above the associated patch of cortex. In contrast, MEG reveals an associated magnetic sink and source response at scalp locations on either side of that scalp location. Because the cortex is crumpled, the variable geometry of different cortical locations (i.e., in a fold versus adjacent to the scalp) means that some locations are more easily seen with EEG whereas others are more easily seen with MEG. Besides providing this different topographic view of brain activity, MEG is a superior technique because magnetic fields are unaffected by the intervening cerebral spinal fluid, skull, skin, hair, etc. In contrast, this intervening material serves to progressively diffuse the voltage potential measured with EEG. Therefore, the voltage potential distributions seen with EEG tend to be spatially large, with opposite values on opposite sides of the head (or even missing in one direction if the other side is below the head). In comparison, the distribution of magnetic flux seen with MEG is more spatially localized, typically revealing both sinks and sources in close proximity on the same side of the head. For this reason, MEG is the measure of choice for ascertaining the reliability of the individual differences in the spatial response.

Although MEG is superior to EEG for its spatial specificity, unique problems are associated with this level of specificity. In particular, it is difficult to perform analyses across participants because a sensor that reveals a sink for one person is just as likely to reveal a source for someone else because (a) the distribution of sinks and sources associated with a current dipole are localized and therefore more sensitive to individual differences in cortical geometry and (b) the MEG sensors are not placed in the same position for each individual, and, therefore, the spatial layout may be affected by the manner in which people position themselves within the MEG chamber. For this reason, we primarily use MEG to assess individual differences, although we developed a geometric localizer technique to reduce the MEG results to a magnitude value, thereby replicating the effects of Experiment 2 with inferential statistics across participants (Tian & Huber, 2008). Due to the problem of head position, we elected not to use dipole modeling with the MEG data because the results of such modeling are complicated by the need to assess where a subject's head is located within the MEG chamber.

Method

Participants

Experiment 3 recruited 18 paid participants, although only 12 of them were able to sufficiently refrain from blinking (the MEG chamber was very dark and there was a tendency for the briefly flashed words to elicit a startle response). Of these 12 participants, 10 of them were used in the reported results. One participant was eliminated due to malfunction of MEG while recording and another was eliminated because of a highly unusual response in one of the conditions as determined through the z scores of the repetition effects (more than 2.5 standard deviations greater than the repetition effects of other participants). Mostly likely this was due to some sort of artifact in one of the conditions, such as movement or blinking. All participants were right-handed native English speakers with no history of neurological disorder.

Experimental Design

The experimental design was identical to Experiment 2 except that every trial was followed by feedback, and the feedback remained onscreen until the participant was ready for the next trial (self-paced). In order to check head position within the MEG, 5 electromagnetic coils were attached to the head of participants before the MEG recording. The locations of the coils were determined with respect to anatomical landmarks (nasion, left and right preauricular points) on the scalp using three-dimensional digitizer software (Source Signal Imaging, Inc. San Diego, CA) and digitizing hardware (Polhemus, Inc., Colchester, VT). The coils were localized with respect to the MEG sensors, both at the beginning and end of the experiment.

During the experiment, participants were supine inside a magnetically shielded room. Before the priming experiment, they listened to 200 repetitions of 250HZ and 1 kHz, 50 ms sinusoidal tone (interstimulus interval randomized between 750 and 1550 ms), with 100 repetitions for each frequency. Auditory-evoked responses to the onset of the pure tones were examined, and the auditory M100 was identified. The auditory M100 is a prominent and robust (across listeners and stimuli) deflection at 100 ms after onset and has been the most investigated auditory MEG response (for review see T. P. Roberts, Ferrari, Stufflebeam, & Poeppel,

2000). A dipole-like pattern (i.e., a source and sink pair) in the magnetic topographic map distributed over the temporal region of each hemisphere was identified for each participant. These responses were used to verify that the participant was positioned properly in the MEG.

MEG Methods

Magnetic signals were recorded using a 160-channel (157 data channels and 3 reference channels), whole-head, axial gradiometer system (KIT, Kanazawa, Japan). The MEG data were acquired with a sampling rate of 500 Hz, filtered online between 1 Hz and 200 Hz, with a notch at 60 Hz (Chait, Poeppel, & Simon, 2006; Liu, Harris, & Kanwisher, 2002). Raw data were noise-reduced using the continuously adjusted least-squares method (CALM; (Adachi, Shimogawara, Higuchi, Haruta, & Ochiai, 2001). A 600-ms time period (100 ms preonset) that was time locked to the first presented prime word was extracted and averaged over the 400 prime responses. These were used to assess individual differences in the M170 response to a singly presented word. A second 600-ms time period that was time locked to the briefly flashed target was extracted and averaged over the 200 novel targets, the 100 targets that repeated the long-duration prime, and the 100 targets that repeated the short-duration prime. Trials with amplitudes > 3pT (~5%) were considered artifactual and discarded. The averages were low pass filtered with a cutoff of 20 Hz (Luck, 2005).

Results and Discussion

After reporting the behavioral results, we report the MEG results in two sections. The first examines MEG responses to the first prime word in order to assess individual differences in the temporal and spatial pattern of the M100 and M170 to a single word. The second examines the MEG responses to the briefly flashed target word in terms of repetition effects. This second analysis uses the correlation (projection) between the target response and prime response to provide a measure of response magnitude that normalizes against the observed individual differences and partially eliminates contamination from overlapping responses from the second prime word. Evidence that this technique improves the reliability of the data averaged over participants is presented below by comparing the raw data based on sensor selection versus this projection technique. In a related paper, Tian and Huber (2008) developed the correlation measures of projection and pattern similarity, demonstrating their effectiveness with the MEG results reported here. That paper also presented spatial topographies of all 10 participants and MEG waveforms for all 157 sensors.

Behavioral Results

The average target flash duration was 92 ms, ranging from 47 ms to 165 ms across individuals. Figure 4 shows the average probability correct for the four conditions. There was a main effect of priming condition, $F(1, 17) = 5.39$, $p < .05$, and a significant interaction between priming condition and prime duration, $F(1, 17) = 21.95$, $p < .001$. For the short prime duration, performance in the target condition was better than performance in the foil

condition, $t(17) = 3.46$, $p < .005$, but for the long prime duration there was no difference, $t(17) = 0.81$, $p = .43$.

MEG Individual Differences

The next two sections test whether there are reliable (significant) differences in the timing of MEG waveform components between individuals and in the topographic patterns associated with MEG waveforms between individuals. Testing the reliability of individual differences informs the question of whether variability in the dipole positions for the fits to Experiments 1 and 2 was due to sampling noise or whether it reflected real differences. As seen in Figure 1, the first prime word was seen on the screen for 1,850 ms in isolation as a single word in the upper location without distraction. Therefore, individual differences in both the timing (temporal differences) and the topography (spatial differences) of the M100 and M170 were assessed by examining the MEG response to the first prime word. This provides a relatively uncontaminated measure of individual differences in the response to visually presented words. Figure 9 shows the results from three participants, demonstrating both different spatial patterns and timing of peak responses (latency) for the M100 and M170 to visually presented words. Because all conditions began with a single word, these data were averaged over all 400 trials for each participant, and so these are fairly reliable responses.

Temporal differences.   Latencies were determined through the root-mean-square across all 157 sensors. The latencies of the first and second major peak after onset were defined as the M100 and M170. The average latency of the M100 was 101.2 ms and ranged from 67 ms to 139 ms (standard deviation of 25.8 ms) and the average latency of the M170 was 174.8 ms and ranged from 135 ms to 203 ms (standard deviation of 23.5 ms). In order to statistically assess the reliability of these differences, the experiment was separated into two halves, and the M100 and M170 latency for each half were determined separately to provide a null hypothesis measure of latency variability. Latency differences (unsigned) between the two halves of the experiment were separately determined within subject (null hypothesis) as well as between subjects (treatment of subject differences). This is similar to the logic of a variance breakdown in an ANOVA, although, in this case, the comparison is made with a simple independent samples t test (unequal variance) comparing the within-subjects first half/second half latency differences versus the between-subjects first half/second half latency differences. Across 10 participants, for both the M100 and the M170, the null hypothesis within latency yielded 10 comparisons (one for each individual), whereas the alternative hypothesis between latency yielded 45 comparisons (45 ways to choose 2 of the 10 individuals). The latency differences for the between-subjects comparison were significantly greater than the within-subjects comparison both for the M100, $t(38.2) = 4.672$, $p < .001$, and the M170, $t(42.6) = 5.419$, $p < .001$. For these independent sample t tests, the equal variance assumption was violated, and so the degrees of freedom was adjusted, resulting in noninteger values.

Spatial differences.   In order to test spatial differences statistically, the spatial pattern of the M100 and M170 (at the appropriate latency) for different individuals was compared
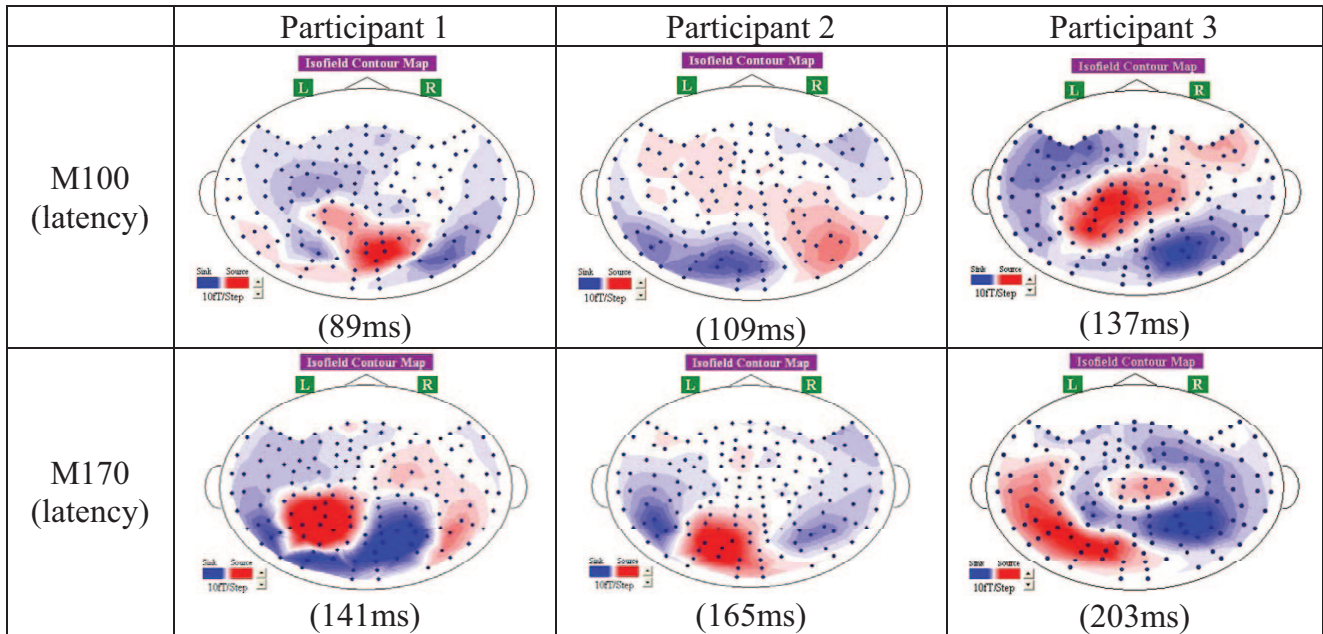
	Participant 1	Participant 2	Participant 3

Figure 9. Experiment 3 latencies and spatial patterns of the M100 and M170 repetition effects for three individuals chosen to highlight individual differences both in the spatial topography and in the timing of the waveforms. Peak latency was determined from the root mean square response to the first prime word, which provided the uncontaminated standard response against which the projection values in Table 3 were calculated.

through a geometric measure of angle in the 157 dimensional sensor space. The similarity in the pattern of two different M100s or M170s was measured through the cosine of the angle between the patterns. Tian and Huber (2008) provided a full description of this angle measure, as well as the projection measure reported next, and statistically compared the results to traditional analyses based on sensor selection (Figure 10 shows this comparison between Part A versus Part B). The angle measure has the advantage that it is unaffected by the magnitude of the response. For example, if two people have identical response patterns but the response for one of the two individuals is twice as large, this comparison technique would reveal that the angle between them was 0, and, thus, the cosine of the angle would be 1.0 (i.e., perfectly similar). Equation 2 shows the vector notation for the cosine of this angle, θ, when comparing two vectors, $\overrightarrow{A}$ and $\overrightarrow{B}$, with these vectors corresponding to the magnetic fields of two different M100s or M170s.

$$\cos\theta = \frac{\overrightarrow{A} * \overrightarrow{B}}{|\overrightarrow{A}||\overrightarrow{B}|} \quad (2)$$

As with the latency differences, the reliability of these spatial similarity differences were determined by breaking the data into first half versus second half of the experiment and this angle was calculated for the 10 within-subjects comparisons versus the 45 between-subjects comparisons. For this comparison it was determined that an equal variance t test was acceptable. The cosine angle for the between-subjects comparison was significantly lower (more dissimilar) than the within-subjects comparison for both the

M100, $t(53) = 4.675$, $p < .001$, and the M170, $t(21.5) = 7.812$, $p < .001$.

Target MEGs and Repetition Effects

Considering these large and reliable individual differences in both the timing and the pattern of the M100s and M170s to visually presented words, we sought to reduce target word responses to a single number that normalized against these individual differences. In the current situation, we determined when the M100 and M170 reached their peak responses, as well as the spatial pattern at that peak, and then calculated a measure of response magnitude at the appropriate time. Not only does this allow a measure that can then be placed into a between-subjects ANOVA, it additionally helps to reduce contamination from the ongoing later components (e.g., M400) in response to the second prime that overlap with the response to the target. The cosine angle test used above assessed similarity between spatial patterns (e.g., whether two M170s were topographically similar); but, for this analysis, we wanted a measure for how much of an M100 or M170 occurred. Therefore, the appropriate measure is the projection of the target word response (at the appropriate time) onto the standard pattern from the prime.

$$|\overrightarrow{B}|\cos\theta = \frac{\overrightarrow{A} * \overrightarrow{B}}{|\overrightarrow{A}|} \quad (3)$$

The left-hand side of Equation 3 is this projection measure in which the vector $\overrightarrow{A}$ is the pattern to the first prime word and vector $\overrightarrow{B}$ is the
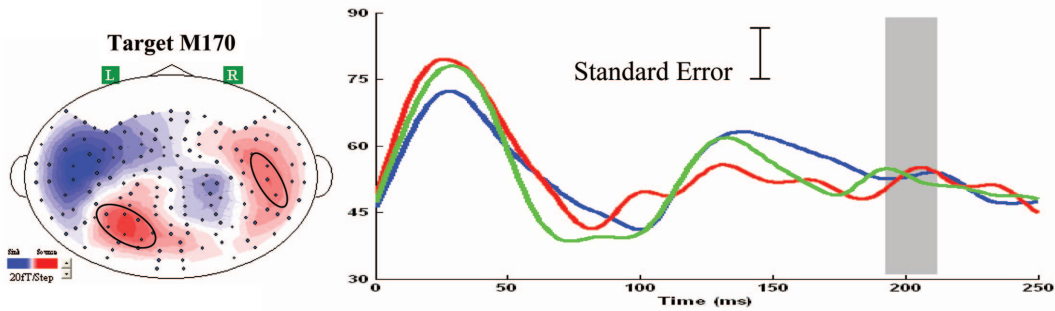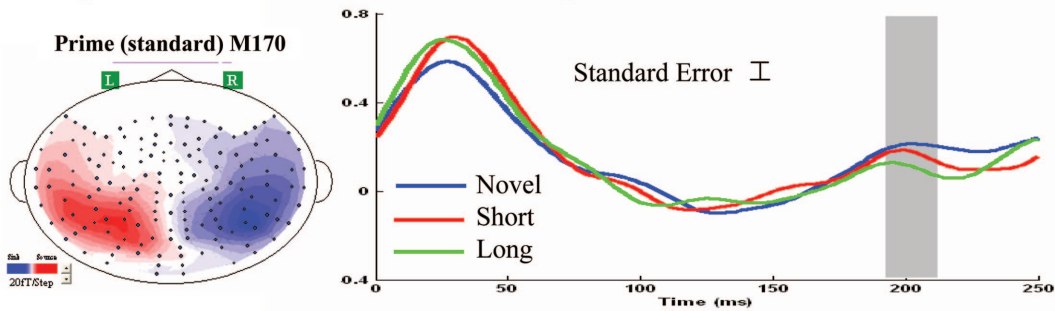
A: Results with Traditional Sensor Selection

B: Projection Results for Multivariate M170 Magnitude

Figure 10. Experiment 3 subject-averaged M170 waveform data for traditional sensor selection (Part A) versus the projection technique (Part B). The graphs show time since the target flash on the x-axis. The legend refers to the prime conditions. As indicated by the standard error bars, the projection results produced more reliable results by normalizing each individual's multivariate magnetoencephalography (MEG) response against their standard M170 response as determined from the first prime. The spatial map in Part A shows the average M170 to the target, which is quite different from the average M170 to the first prime shown in Part B, presumably due to the overlap between the response to the target and the responses to the second prime word. The gray bars of the graphs highlight the time window where the M170 to the target should have occurred, although the magnitude of response was small due to the short duration of the masked target and overlap with the response to the second prime word. For comparison, note the large M170 seen in both graphs between 0 and 50 ms, which is due to the second prime word presented 150 ms prior to the zero point when the target was presented. The graph of Part A shows the three conditions as calculated from the 10 circled sensors with the largest source response. In contrast to these results, the graph of Part B shows the projection results, which revealed a small M170 to the target and also revealed that the M170 was less in the long condition but not the short condition.

pattern to the target word. These vectors were determined by averaging over a 22-ms time window centered on the individual M100 or M170 peak responses. Thus, each participant yielded three data points for the M100 and M170 for each of the three priming conditions. This is analogous to a localizer task in functional magnetic resonance imagining (fMRI; e.g., Kanwisher, Tong, & Nakayama, 1998) that first determines which region of the brain is of interest separately for each participant (here based on the MEG response to the prime), and then examines changes in the response of the appropriate region across the conditions (i.e., the MEG response to targets in the repeated vs. novel conditions).

Before normalizing with this geometric projection, a critical first step involves comparing the three target conditions for their spatial similarity using the cosine angle test. If the three conditions are found to be spatially dissimilar in their response, then any apparent magnitude difference from the projection onto the standard prime response may be due to spatial similarity differences between the conditions. Conversely, if the

three conditions are found to be spatially similar, then their projection onto the standard prime response provides a relatively assumption-free indication that the change was in the magnitude of the cortical response rather than in the distribution of the cortical response. Statistical reliability was again determined by dividing the experiment into trials from the first half of the experiment versus trials from the second half of the experiment, with the key comparison of within-condition versus between-condition angle tests (these were always calculated within subject). Because there were three conditions, this entailed three different possible comparisons between conditions, which were then averaged in comparison to the average of the three within-condition comparisons, yielding one between number and one within number for each participant. In a dependent samples test across the 10 participants, there was no significant difference in the similarity (angle test) of the between-conditions comparisons versus the within-condition comparisons, and this was true both for the M100, $t(9) = 0.705$, $p =$

.499 and the M170, $t(9) = -1.216$, $p = .255$. Thus, the topographic patterns across the target responses in the three conditions were similar enough to allow subsequent magnitude comparisons with the projection test.[2] This also implies that the same distribution of cortical responses was involved in the three target conditions such that any magnitude differences are real and not due to differential recruitment of cortical areas in some conditions but not others.

Figure 10 demonstrates that the projection measure can produce more reliable results across participants by showing the average MEG waveform for the M170 with a traditional method of sensor selection (Part A) versus the project method (Part B). The topographic maps show the standard M170 response to the first prime (Part B) versus the M170 response to the brief target (Part A), which are very different, presumably due to overlapping responses to the second prime or the M100 to the mask. Because the projection measure normalizes against individual differences and partially removes overlapping responses, the error bar for the graph in Part B is relatively smaller compared to Part A (i.e., the error bar is small relative to the differences between conditions seen during the time window of the target M170 as highlighted by the gray bar). Furthermore, the graph in Part B reveals a small positive deflection during the time window of the M170 for all conditions, whereas the M170 for the target is not apparent in Part A, suggesting that projection has recovered the M170 signal. Providing a more reliable measure, Table 3 reports the projection values averaged over the M170 time window.

Comparing across the three priming conditions in terms of the projection measure of magnitude (see Table 3), there were no differences for the M100, $F(2, 18) = 1.383$, $p = .276$, but there were for the M170, $F(2, 18) = 5.202$, $p < .025$. The M170 to a repeated word was smaller than the novel condition when the prime was presented for the long duration, $t(9) = 3.916$, $p < .01$, but there was no repetition priming effect for the short duration, $t(9) = 0.919$, $p = .382$. In conclusion, Experiment 2 found that there were no P100 ($\sim$100 ms) repetition effects but there was less of an N170 ($\sim$170 ms) when the target followed a long-duration prime and this pattern is replicated in terms of the M100 and M170 using the same two-prime-word paradigm as Experiment 2.

General Discussion

With threshold identification of words and immediate repetition priming, it has been observed that short prime durations produce positive priming but long prime durations eliminate this effect or even produce negative priming (Huber, 2008; Huber, Shiffrin,

Lyle, et al., 2002; Huber et al., 2001; Huber, Shiffrin, Quach, et al., 2002; Weidemann et al., 2005, 2008). In order to differentiate between a perceptual basis versus a decision basis of this prime duration effect, we examined early perceptual brain responses to the target word, which was presented prior to the choice words and initiation of the decision process. Supporting the claim that these effects may have a perceptual basis, we found that ERP repetition effects were affected by prime duration and replicated these results with MEG.

The theory that guides this work assumes that perceptual responses are integrated over time, potentially causing source confusions (short primes), but that with sufficient processing (long primes), this source confusion is minimized (i.e., prime and target are separated) through habituation to previously identified items. This theory has successfully explained many behavioral priming results in which priming effects are reduced, eliminated, or even reversed by supraliminal primes (Huber, 2008; Huber, Clark, Curran, & Winkielman, in press; Huber & Cousineau, 2004; Huber & O'Reilly, 2003; Rieth & Huber, 2005, 2008). Using parameters that previously specified behavioral results as a function of prime duration, we used this neural habituation model to produce a priori predictions for the direction of P100 and N170 repetition effects and the role of prime duration by assuming that letter processing underlies the P100 and word processing underlies the N170. Our results are significant both qualitatively and quantitatively. Qualitatively, we found clear evidence of perceptual ERP and MEG effects with immediate repetition. Furthermore, these effects were modulated by prime duration. When only a single prime word was presented (Experiment 1), we observed less of a response for the P100 in the repeated condition (less positive), but this repetition effect was smaller following a long-duration prime. Likewise, the N170 component in the repeated condition was larger (more negative) following a short-duration prime and this effect also interacted with prime duration, such that the N170 was smaller (less negative) following a long-duration prime. Experiments 2 and 3 used two different prime words on every trial in order to control for attentional factors, and both experiments found that the P100/ M100 repetition effects were eliminated, but there were still repetition effects for the N170/M170, with smaller responses for repeated words only following long-duration primes. The difference between Experiment 1 versus Experiments 2 and 3 was expected due to the extra competition that takes place within the word identification system with the use of two different primes rather than a single prime word. The overall pattern of results, and the interactions with prime duration in particular, provides converging evidence that the corresponding behavioral priming effects

Table 3

Target Word Responses Projected Onto Prime Responses To Normalize Against Individual Differences

Projection	M100		M170	
	Mean	SE	Mean	SE
Novel	.439	.147	.207	.104
Short	.422	.180	.176	.109
Long	.329	.168	.113**	.113

Note. ** MEG repetition effects significant at the .01 level.

[2] The projection test is useful for normalizing against individual differences even though the angle test conclusion that the patterns are sufficiently similar relies upon accepting the null hypothesis, with some unknown type II error rate. In other words, use of the projection test for assessing statistically reliable differences does not rely on accepting the null hypothesis for the angle test, but interpretation of the differences in terms of magnitude does rely on this acceptance. As shown in Figure 10 and Table 3, the projection measure produced reliable and significant results, whereas the angle test failed to produce significant results. This indicates that an interpretation of the differences between conditions in terms of magnitude is more likely than an interpretation in terms of topographic similarity.

and the interaction with prime duration is due to relatively auto-matic perceptual processes, such as is assumed by habituation model of Huber and O'Reilly (2003), rather than strategic respond-ing and rather than a decisional process specifically in response the choice words.

These results are quantitatively important because the simulated neural network accounted for the moment-by-moment ERP repe-tition effects with only dipole positions and static dipole vector moments as free parameters (i.e., the neural network specified the neural dynamics in advance based on the behavioral data). Our interpretation is novel and assumes that the neural sources of ERP components are isolated neither in time nor location, with the ongoing responses of early sources contributing to the observed effects of later ERP components. The prediction of the model was for less of a cortical response to a repetition due to integration across multiple items (i.e., the activation of a brain region to the presentation of two different words is greater than repeating the same word). Furthermore, the magnitude of this repetition effect was predicted to be smaller following a long-duration prime due to separation. The observed P100 effects (Experiment 1) were di-rectly in accord with this prediction. The N170 effects were also in agreement with this account assuming that the source dipoles were oppositely oriented (see Figure 3). In this manner, the observation of larger N170s (Experiment 1) or a lack of an N170 effect (Experiments 2 and 3) following a short-duration prime was ex-plained as the ongoing source of the P100/M100 effect partially or fully counterbalancing the N170/M170 effect. Finally, the emer-gence of decreased N170/M170s following long-duration primes was explained as the unmasking of the "pure" underlying N170/M170 repetition effect because the P100/M100 repetition effects were diminished in size following long-duration primes due to habituation.

Experiment 3 used MEG, which allowed better spatial localiza-tion for the cortical source underlying an electrophysiological response. Besides replicating Experiment 2 with a different tech-nique, this increased spatial specificity allowed a test for the reliability of individual differences. The M100 and M170 re-sponses to visually presented words were found to be highly variable across individuals in terms of both the latency of the responses and the pattern of the responses across the sensors. Assuming that anatomical structures of different individuals do not differ as greatly as might be implied by the differences in the electromagnetic flux patterns (see Figure 9), it seems likely that these patterns reflect the combined response across a distributed array of cortical responses, with the individual differences arising from different mixtures of the array rather than differences in a single source. Therefore, although the dipole fits were useful for testing the model's dynamics, the dipole positions and orientations should not be taken in a literal manner and instead are best thought of as capturing the central tendency for the distribution of cortical responses, with this central tendency differing across individuals. In order to normalize against these individual differences to di-rectly test the model's predictions in terms of magnitude, Exper-iment 3 used the prime's response as a standard latency and pattern.

A question that often arises in the use of a computational model for explaining data is whether the model is too flexible and what types of data cannot be addresses by the model (Pitt et al., 2003; Pitt & Myung, 2002; Pitt et al., 2002; S. Roberts & Pashler, 2000).

Our finding of early repetition effects in response to the target, prior to the choice words, strongly suggests a perceptual locus of these priming effects, but these results cannot unequivocally iden-tify habituation as the root cause. Therefore, the issue of model flexibility is important in assessing the strength of an explanation in terms of habituation, and we addressed this issue in several different ways. First, we used previously published parameters based on behavioral data to produce predictions with the assump-tion of exactly equal and oppositely oriented responses for the letter and word levels of processing. In reporting the results, we converted the mathematical sign of all electrodes in relation to posterior scalp regions so that the obtained results (Figure 7) could be directly compared to these original a priori predictions (Figure 3B). Second, we derived a new method for statistically assessing the goodness of fit to ERP data. Thus, a chi-square value was calculated through the likelihood ratio for the probability of the data under the model, versus the probability of the data from a model that perfectly fit the obtained results. To obtain these probabilities, we assumed normally distributed error variance in sampling ERPs and used measures of ERP variance for each individual to calculate a goodness of fit for each individual. These chi-square values revealed that the model produced a valid fit for all individuals in accounting for 9,288 data points of an individual with just 15 free parameters. Third, we engaged in model com-parison by also fitting the popular BESA dipole methodology (Scherg, 1990), which assumes that equivalent dipoles are inde-pendent across time steps and can take on any value regardless of other time steps. This increased the number of parameters by a factor of 10, but this model did not fare significantly better for the group-averaged data. Thus, the a priori dynamics of neural habit-uation based on group-averaged behavioral data were supported.

Comparison with Other Studies

In the study of memory and language processing, there is a sizable literature examining ERPs to immediately repeated words both within and across modality (e.g., Bentin & McCarthy, 1994; Doyle, Rugg, & Wells, 1996; Rugg, Mark, Gilchrist, & Roberts, 1997; Rugg & Nieto-Vegas, 1999), but very few of these studies examined immediate repetition effects on early (< 200 ms) ERP components. However, Holcomb, Grainger, and colleagues de-scribed a P150 effect that is smaller when a word is preceded by a masked prime of the same word or letter compared to a different word (Holcomb & Grainger, 2006) or letter (Petit, Midgley, Holcomb, & Grainger, 2006). For the word-priming study (Hol-comb & Grainger, 2006), it is interesting to note that the P150 seems to be superimposed on a large negative-going N170 com-ponent. Thus, the effect could be described as a P150 attenuation or an N170 enhancement, as noted by Holcomb and Grainger (2006). According to our model, any effects of repetition attenuate neural activity, suggesting that these results represent the summa-tion of oppositely oriented dipoles (i.e., attenuation of a positive dipole producing the repetition effect combined with a larger negative dipole producing the negativity).

In contrast to word studies, early ERP and MEG immediate repetition effects with face and object processing have been stud-ied more extensively. The most consistent finding is for smaller N170/M170s following immediately repeated compared to nonre-peated faces or objects as measured both with ERP (Campanella et

al., 2000; Heisz, Watter, & Shedden, 2006; Henson, Rylands, Ross, Vuilleumeir, & Rugg, 2004; Itier & Taylor, 2004; Jemel, Pisani, Calabria, Crommelinck, & Bruyer, 2003; Kovacs et al., 2006; Martens, Schweinberger, Kiefer, & Burton, 2006) and MEG (Harris & Nakayama, 2007; Ishai, Bikle, & Ungerleider, 2006). These results are consistent with our model, which predicts less cortical response to immediate repetitions. However, several other studies have failed to find effects of immediate face or object repetitions on early ERP (Henson et al., 2003; Schweinberger, Pickering, Burton, & Kaufmann, 2002; Werheid, Alpay, Jentzsch, & Sommer, 2005) or MEG (Halgren, Raij, Marinkovic, Jousmaki, & Hari, 2000; Penney, Maess, Busch, Derrfuss, & Mecklinger, 2003) responses. Such inconsistent effects are also understandable in light of our results and theory, considering that repetition effects observed at the scalp can be influenced by several factors, including prime duration and the spatiotemporal summation of multiple cortical sources that may have counteracting influences because of oppositely oriented dipoles. We do not claim that the same neural substrates underlie object, face, and word identification, but rather that the cortical dynamics should be similar. As such, we expected and recently observed analogous behavioral priming effects for face stimuli (Rieth & Huber, 2008).

In modeling our results, it was not necessary to specify what type of information Levels 2 and 3 of the neural network encode, considering that we did not manipulate orthographic or semantic similarity. However, by assuming Level 2 encodes orthographic information (letter processing) and Level 3 encodes lexical-semantic information (word processing), Huber (2008) accounted for behavioral experiments that manipulated the degree of orthographic and semantic similarity between primes and primed choice words. We do not claim that P100s only occur with letter processing or that N170s only occur with word processing. Instead, in keeping with work examining faces, nonwords, and other visual objects, our claim is that the P100 is sensitive to processing of component features (e.g., letters or face parts) and that the N170 is sensitive to the processing of whole objects (e.g., words or faces). The model only represents word processing, although similar object modules presumably exist for faces and other visual objects, with parts and wholes producing P100 and N170 effects. However, these responses to other objects are perhaps of different magnitudes and topography, reflecting different cortical regions within the ventral stream of visual processing (i.e., the 'what' pathway). From this larger representational perspective, it is important that we restricted our analyses to repetition effects (repeated minus novel) in order to focus the results on aspects unique to letters and words.

Consistent with these representational assumptions, Sereno, Rayner, and Posner (1998), observed orthographic effects for the P100 (pseudowords versus consonant strings) and word frequency effects for the N170. Holcomb and Grainger have also related the P150 with orthographic processing based on similar repetition effects for words and letters (Grainger & Holcomb, in press; Holcomb & Grainger, 2006; Petit et al., 2006). More recently, Sereno, Brewer and O'Donnell (2003) found that the N170 response to a semantically ambiguous word was affected by whether the sentence context provided disambiguating information. This result supports the claim that Level 3 of the model, which produced N170 repetition effects, is involved in the early stages of lexical-semantic processing.

Immediate repetition effects are difficult to measure with fMRI, considering the relatively poor temporal resolution of that technique. Nonetheless, Dehaene et al. (2001) found a case-independent "repetition suppression" effect in the left fusiform gyrus for immediately repeated words (our interpretation would call this an "integration effect," such that two words provide greater activation than repeating a single word). In a follow-up experiment using event-related fMRI, Dehaene et al. (2002) found that this region responded to written but not spoken words, and so they termed this a visual word form area. In reconciling our data with these results and with the results of Sereno and colleagues, it could be that Level 3 of the network is analogous to Dehaene's visual word form area; but, nonetheless, the representation in this visual area may be structured according to lexical-semantic similarity.

Conclusions

It may seem counterintuitive that the way to separate items in time is through habituation, but this separation is designed to allow accurate identification of other, unrelated items (novel condition) at the cost of identical or highly similar items (repeated condition). Viewed in this way, this mechanism is sensible provided that things do not often immediately repeat. This temporary habituation serves to reduce source confusion between previously identified items and subsequent items. In the absence of this habituation, things blend together across presentations resulting in the tendency to mistake the prime for the target. However, once prime and target are separated through habituation, these effects are reversed, producing a tendency to perceive anything but the prime, which lowers performance in the repeated condition.

Qualitatively, this theory predicted smaller neural responses for repeated words, with these effects occurring in early perceptual responses to the target word, prior to the decision process. Furthermore, this theory predicted that these repetition effects would modulate with prime duration. Using two different electrophysiological recording techniques, we confirmed these qualitative predictions.

The basic dynamics of the model specify that brief prime presentations enhance target-specific responses but longer prime presentations result in smaller target-specific responses. However, this prediction does not directly map into electrophysiological measures because scalp recordings are sensitive to the entire word form area and reflect the combined activation of prime and target. If source confusion between prime and target is due to the simultaneous perceptual activation of prime and target, this implies that the total activation is largest when prime and target differ in identity. Therefore, both source confusion and habituation were expected to produce repetition deficits: Source confusion would produce an apparent deficit because two different simultaneously active words (i.e., a prime followed by a different target) would sum to a larger response than a repetition of a single word, and habituation would also cause a deficit through a direct reduction in the target response. In sum, the general prediction was that perceptual responses should be smaller with repetitions as measured with scalp recordings and, furthermore, that these effects should modulate with prime duration in moving from integration to separation.

These experiments identified two early electrophysiological responses that were affected by the repetition of visually presented words. Furthermore, these responses were modulated by prime duration, with the P100 effect decreasing with increased prime duration, whereas N170/M170 effect either reversed or went from a null effect to a repetition decrease with increased prime duration. This provides converging evidence for the claim that behavioral priming in this paradigm results from changes in perceptual processing. As expected from this perceptual account, we observed these electrophysiological effects prior to presentation of the choice words. We cannot rule out the possibility of strategic decision, but we appeal to parsimony in presenting a model that simultaneously accounted for both ERP and behavioral data without reference to any postperceptual factors. Furthermore, this was accomplished in an a priori manner, using the published network dynamics.

Our interpretation of the electrophysiological data is novel, proposing that different stages of perceptual processing combine both in time and space to produce scalp potentials and magnetic fields. This highlights the potential hazards of assuming different electrophysiological components uniquely relate to specific isolated processing events. Instead, we champion an integrative dynamic methodology for interpretation of electrophysiological data.

References

Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science, 275*(5297), 220–224.

Adachi, Y., Shimogawara, M., Higuchi, M., Haruta, Y., & Ochiai, M. (2001). Reduction of nonperiodic environmental magnetic noise in MEG measurement by continuously adjusted least squares method. *IEEE Transactions on Applied Superconductivity, 11*(1), 669–672.

Baillet, S., & Garnero, L. (1997). A Bayesian approach to introducing anatomo-functional priors in the EEG/MEG inverse problem. *IEEE Transactions on Biomedical Engineering, 44*(5), 374–385.

Batchelder, W. H., & Riefer, D. M. (1990). Multinomial Processing Models of Source Monitoring. *Psychological Review, 97*(4), 548–564.

Bekker, E. M., Kenemans, J. L., & Verbaten, M. N. (2005). Source analysis of the N2 in a cued Go/NoGo task. *Cognitive Brain Research, 22*(2), 221–231.

Bentin, S., & McCarthy, G. (1994). The effects of immediate stimulus repetition on reaction time and event-related potentials in tasks of different complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(1), 130–149.

Berg, P., & Scherg, M. (1994). A fast method for forward computation of multiple-shell spherical head models. *Electroencephalography and Clinical Neurophysiology, 90*(1), 58–64.

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*(6), 796–806.

Chait, M., Poeppel, D., & Simon, J. Z. (2006). Neural response correlates of detection of monaurally and binaurally created pitches in humans. *Cerebral Cortex, 16*(6), 835–848.

Collins, L. M., Fidler, P. L., Wugalter, S. E., & Long, J. D. (1993). Goodness-of-Fit Testing for Latent Class Models. *Multivariate Behavioral Research, 28*(3), 375–389.

Dehaene, S., Le Clec, H. G., Poline, J. B., Le Bihan, D., & Cohen, L. (2002). The visual word form area: A prelexical representation of visual words in the fusiform gyrus. *Neuroreport, 13*(3), 321–325.

Dehaene, S., Naccache, L., Cohen, L., Bihan, D. L., Mangin, J. F., Poline, J. B., et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nature neuroscience, 4*(7), 752–758.

Dien, J. (1998). Issues in the application of the average reference: Review, critiques, and recommendations. *Behavior Research Methods, Instruments, and Computers, 30,* 34–43.

Doyle, M. C., Rugg, M. D., & Wells, T. (1996). A comparison of the electrophysiological effects of formal and repetition priming. *Psychophysiology, 33*(2), 132–147.

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap. Monographs on statistics and applied probability, Vol. 57.* New York: Chapman & Hall/CRC.

Eimer, M. (1995). Stimulus-response compatibility and automatic response activation–Evidence from psychophysiological studies. *Journal of Experimental Psychology: Human Perception and Performance, 21*(4), 837–854.

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(4), 325–350.

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*(4), 669–687.

Gonzalez, R., & Griffin, D. (2001). Testing parameters in structural equation modeling: Every "one" matters. *Psychological Methods, 6*(3), 258–269.

Grainger, J., & Holcomb, P. J. (in press). Neural constraints on a functional architecture for word recognition. In P. Cornelissen, P. Hansen, M. Kringlebach, & K. Pugh (Eds.), *The neural basis of reading.* Oxford, England: Oxford University Press.

Halgren, E., Baudena, P., Clarke, J. M., Heit, G., Liegeois, C., Chauvel, P., et al. (1995). Intracerebral Potentials to Rare Target and Distracter Auditory and Visual-Stimuli. 1. Superior Temporal Plane and Parietal Lobe. *Electroencephalography and Clinical Neurophysiology, 94*(3), 191–220.

Halgren, E., Raij, T., Marinkovic, K., Jousmaki, V., & Hari, R. (2000). Cognitive response profile of the human fusiform face area as determined by MEG. *Cerebral Cortex, 10*(1), 69–81.

Hamalainen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., & Lounasmaa, O. V. (1993). Magnetoencephalography: Theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics, 65*(2), 413–497.

Handy, T. C. (2005). *Event-related potentials: A methods handbook.* Cambridge, MA: MIT Press.

Harris, A., & Nakayama, K. (2007). Rapid face-selective adaptation of an early extrastriate component in MEG. *Cerebral Cortex, 17*(1), 63–70.

Heisz, J. J., Watter, S., & Shedden, J. M. (2006). Progressive N170 habituation to unattended repeated faces. *Vision Research, 46*(1–2), 47–56.

Henson, R. N., Goshen-Gottstein, Y., Ganel, T., Otten, L. J., Quayle, A., & Rugg, M. D. (2003). Electrophysiological and haemodynamic correlates of face perception, recognition and priming. *Cerebral Cortex, 13*(7), 793–805.

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*(4), 1674–1689.

Hochhaus, L., & Johnston, J. C. (1996). Perceptual repetition blindness effects. *Journal of Experimental Psychology: Human Perception and Performance, 22*(2), 355–366.

Hoel, P. G. (1971). *Introduction to mathematical statistics* (4th ed.). New York: Wiley.

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*(10), 1631–1643.

Hopf, J. M., Vogel, E., Woodman, G., Heinze, H. J., & Luck, S. J. (2002). Localizing visual discrimination processes in time and space. *Journal of Neurophysiology, 88*(4), 2088–2095.

Huber, D. E. (2006). Computer simulations of the ROUSE model: An analytic method and a generally applicable techniques for producing parameter confidence intervals. *Behavior Research Methods, 38,* 557–568.

Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General, 137*(2), 324–347.

Huber, D. E., Clark, T. F., Curran, T., & Winkielman, P. (in press). Effects of repetition priming on recognition memory: Testing a perceptual fluency-disfluency model. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

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

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, 27*(3), 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*(6), 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*(1), 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*(5), 745–757.

Humphreys, G. W., Besner, D., & Quinlan, P. T. (1988). Event perception and the word repetition effect. *Journal of Experimental Psychology: General, 117*(1), 51–67.

Inui, K., & Kakigi, R. (2006). Temporal analysis of the flow from V1 to the extrastriate cortex in humans. *Journal of Neurophysiology, 96*(2), 775–784.

Ishai, A., Bikle, P. C., & Ungerleider, L. G. (2006). Temporal dynamics of face repetition suppression. *Brain Research Bulletin, 70*(4–6), 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*(4), 1518–1532.

Jacoby, L. L., & Whitehouse, K. (1989). An illusion of memory: False recognition influenced by unconscious perception. *Journal of Experimental Psychology: General, 118*(2), 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*(2), 431–446.

Jung, T. P., Makeig, S., Westerfield, M., Townsend, J., Courchesne, E., & Sejnowski, T. J. (2000). Removal of eye activity artifacts from visual event-related potentials in normal and clinical subjects. *Clinical Neurophysiology, 111*(10), 1745–1758.

Junghofer, M., Elbert, T., Tucker, D. M., & Braun, C. (1999). The polar average reference effect: A bias in estimating the head surface integral in EEG recording. *Clinical Neurophysiology, 110*(6), 1149–1155.

Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (1991). *Principles of neural science* (3rd ed.). New York: Elsevier.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition, 27*(2), 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*(5), 1249–1260.

Kanwisher, N. G., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition, 68*(1), B1–B11.

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

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*(5), 742–753.

Kucera, 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 after effects. *Nature Neuroscience, 4*(1), 89–94.

Liu, J., Harris, A., & Kanwisher, N. (2002). Stages of processing in face perception: An MEG study. *Nature Neuroscience, 5*(9), 910–916.

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*(6), 609–616.

Luck, S. J. (2005). *An introduction to the event-related potential technique.* Cambridge, Mass.: MIT Press.

MacCallum, R. C., Browne, M. W., & Cai, L. (2006). Testing differences between nested covariance structure models: Power analysis and null hypotheses. *Psychological Methods, 11*(1), 19–35.

Makeig, S., Jung, T. P., Bell, A. J., Ghahremani, D., & Sejnowski, T. J. (1997). Blind separation of auditory event-related brain responses into independent components. *Proceedings of the National Academy of Sciences, USA, 94*(20), 10979–10984.

Makeig, S., Westerfield, M. J., T.-P., Enghoff, S., Townsend, J., Courchesne, E., et al. (2002). Dynamic brain sources of visual evoked responses. *Science, 295,* 690–694.

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

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.

Masson, M. E. J., & Borowsky, R. (1998). More than meets the eye: Context effects in word identification. *Memory & Cognition, 26*(6), 1245–1269.

McClelland, J. L., & Rumelhart, D. E. (1981). An Interactive activation model of context effects in letter perception. Part 1. An account of basic findings. *Psychological Review, 88*(5), 375–407.

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*(2), 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*(2), 309–321.

Mosher, J. C., Baillet, S., & Leahy, R. M. (1999). EEG source localization and imaging using multiple signal classification approaches. *Journal of Clinical Neurophysiology, 16*(3), 225–238.

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*(5), 723–739.

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

Penney, T. B., Maess, B., Busch, N., Derrfuss, J., & Mecklinger, A. (2003). Cortical activity reduction with stimulus repetition: A whole-head MEG analysis. *Cognitive Brain Research, 16*(2), 226–231.

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*(4), 674–681.

Pitt, M. A., Kim, W., & Myung, I. J. (2003). Flexibility versus generalizability in model selection. *Psychonomic Bulletin & Review, 10*(1), 29–44.

Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Science, 6*(10), 421–425.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting

among computational models of cognition. *Psychological Review, 109*(3), 472–491.

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

Qiu, J., Luo, Y. J., Wang, Q. H., Zhang, F. H., & Zhang, Q. L. (2006). Brain mechanism of Stroop interference effect in Chinese characters. *Brain Research, 1072*(1), 186–193.

Regan, D. (1989). *Human brain electrophysiology: Evoked potentials and evoked magnetic fields in science and medicine.* New York: Elsevier.

Reinvang, I., Magnussen, S., Greenlee, M. W., & Larsson, P. G. (1998). Electrophysiological localization of brain regions involved in perceptual memory. *Experimental Brain Research, 123*(4), 481–484.

Rhodes, G., Robbins, R., Jaquet, E., McKone, E., Jeffery, L., & Clifford, C. W. G. (2005). Adaptation and face perception: How aftereffects implicate norm-based coding of faces. In C. W. G. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world: Adaptation and after-effects in high-level vision* (pp. 213–240). New York: Oxford University Press.

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 E. Associates (Ed.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 1856–1861). Hillsdale, NJ.

Rieth, C. A., & Huber, D. E. (2008). *Priming and habituation for faces: Individual differences and inversion effects.* Manuscript submitted for publication.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107*(2), 358–367.

Roberts, T. P., Ferrari, P., Stufflebeam, S. M., & Poeppel, D. (2000). Latency of the auditory evoked neuromagnetic field components: Stimulus dependence and insights toward perception. *Journal of Clinical Neurophysiology, 17*(2), 114–129.

Rugg, M. D., Mark, R. E., Gilchrist, J., & Roberts, R. C. (1997). ERP repetition effects in indirect and direct tasks: Effects of age and interitem lag. *Psychophysiology, 34*(5), 572–586.

Rugg, M. D., & Nieto-Vegas, M. (1999). Modality-specific effects of immediate word repetition: Electrophysiological evidence. *Neuroreport, 10*(12), 2661–2664.

Samuel, A. G., & Kat, D. (2003). Inhibition of return: A graphical meta-analysis of its time course and an empirical test of its temporal and spatial properties. *Psychonomic Bulletin & Review, 10*(4), 897–906.

Scherg, M. (1990). Fundamentals of dipole source potential analysis. In: Auditory evoked magnetic fields and electric potentials. In F. Grandori, M. Hoke & G. L. Romani (Eds.), *Advances in Audiology* (Vol. 6, pp. 40–69): Karger, Basel.

Schweinberger, S. R., Pickering, E. C., Burton, A. M., & Kaufmann, J. M. (2002). Human brain potential correlates of repetition priming in face and name recognition. *Neuropsychologia, 40*(12), 2057–2073.

Sereno, S. C., Brewer, C. C., & O'Donnell, P. J. (2003). Context effects in word recognition: Evidence for early interactive processing. *Psychol Sci, 14*(4), 328–333.

Sereno, S. C., Rayner, K., & Posner, M. I. (1998). Establishing a time-line of word recognition: Evidence from eye movements and event-related potentials. *Neuroreport, 9*(10), 2195–2200.

Srinivasan, R., Nunez, P. L., Tucker, D. M., Silberstein, R. B., & Cadusch, P. J. (1996). Spatial sampling and filtering of EEG with spline laplacians to estimate cortical potentials. *Brain Topography, 8*(4), 355–366.

Stok, C. J. (1986). *The inverse problem in EEG and MEG with application to visual evoked responses.* The Hague, the Netherlands: CIP Gegevens Koninklijke Bibliotheek.

Thierry, G., Doyon, B., & Demonet, J. F. (1998). ERP mapping in phonological and lexical semantic monitoring tasks: A study complementing previous PET results. *Neuroimage, 8*(4), 391–408.

Tian, X., & Huber, D. E. (2008). Measures of spatial similarity and response magnitude in MEG and scalp EEG. *Brain Topography, 20*(3), 131–141.

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*(2), 719–723.

Tucker, D. M. (1993). Spatial sampling of head electrical fields: The geodesic sensor net. *Electroencephalography and Clinical Neurophysiology, 87*(3), 154–163.

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.

Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2008). Prime diagnosticity in short-term repetition priming: Is primed evidence discounted even when it reliably indicates the correct answer? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*(2), 257–281.

Werheid, K., Alpay, G., Jentzsch, I., & Sommer, W. (2005). Priming emotional facial expressions as evidenced by event-related brain potentials. *International Journal of Psychophysiology, 55*(2), 209–219.

Whittlesea, B. W. A., Dorken, M. D., & Podrouzek, K. W. (1995). Repeated Events in Rapid Lists: Part 1. Encoding and Representation. *Journal of Experimental Psychology: Learning Memory, and Cognition, 21*(6), 1670–1688.

Woldorff, M. G. (1993). Distortion of ERP averages due to overlap from temporally adjacent ERPS: Analysis and correction. *Psychophysiology, 30*(1), 98–119.